ETHICAL PRACTICE AND OBJECTIVE REASONS

Kantian and Wittgensteinian Themes in the Objectivity of Ethical Reasons

D.PHIL. THESIS



JAMES CORNWELL

BRASENOSE COLLEGE

ETHICAL PRACTICE AND OBJECTIVE REASONS

Kantian and Wittgensteinian Themes in the Objectivity of Ethical Reasons

ETHICAL PRACTICE AND OBJECTIVE REASONS:
Kantian and Wittgensteinian Themes in the Objectivity of Ethical Reasons

James Cornwell
Brasenose College

ABSTRACT

D.Phil. Thesis
Trinity 2000

This thesis develops a 'response-dependent' account of the objectivity of ethical judgment. It uses the Kantian device of the Categorical Imperative (CI) as the critical reflective test for convergence in judgment that is genuinely objectivity-sustaining. After a discussion of the challenges to ethical objectivism, it is argued that Wittgenstein's consideration of rule-following provides a compelling case for rejecting attempts to ground objectivity in a manner that is external to normative practices. This removes one potent source of hostility to ethical objectivity, but it does not establish the truth of this position. Within the constraints of a broadly Wittgensteinian conception of objectivity, plausible accounts of ethical objectivity will be 'response-dependent'. The key issue facing such theories is whether they can provide a model of practical rationality that delivers suitably demanding, objectivity-sustaining standards of rational criticism. I examine two approaches to this issue. The first (suggested by John McDowell) takes Wittgenstein's rule-following considerations as implying a virtue-centred approach to ethical deliberation. On this view, virtuous agents' uncodifiable responses determine correct judgment. But the case for this theoretically modest model of objectivity has not been successfully made, and the approach has insufficient resources to yield genuine objectivity. Instead, appeal must be made to some more articulated account of practical rationality. The second, Kantian, model delivers this because the CI test provides a mechanism for objectivity-sustaining reflective criticism of agents' reasons for action, without importing any illicit codification of rationality. The CI can also be used for critical reflection on agents' appreciation of the ethical features of their situation. In this way the model can avoid the problem of empty formalism and offer a weak form of ethical cognitivism. This model also suggests a potential route to plausible approaches to some other problematic issues in ethical theory, such as akrasia and the categorical nature of moral reasons.

CONTENTS

List of Abbreviations v

Acknowledgements vi

CHAPTER 1. INTRODUCTION 1

1. The Challenge 1
2. A Preview of the Project 4
3. A Few Methodological Remarks 7

CHAPTER 2. ETHICAL OBJECTIVITY: WHAT IS AT STAKE? 15

1. Some Terminological Preliminaries 15
2. Routes through the Maze 18
3. Equating Objectivity with Truth-Aptness 20
4. Minimalist Truth and Objectivity 24
5. The Rejection of Objectivism 29
6. Essential Contestability and Objectivity 33
7. Judgmental Convergence and Explanation 39
8. Conclusions 48

CHAPTER 3. RULE-FOLLOWING AND OBJECTIVITY 50

1. Objectivity, Meaning and Rules 50
2. How Not to Follow a Rule (1): Mental States 55
3. How Not to Follow a Rule (2): Dispositional States 58
4. How Not to Follow a Rule (3): Platonism 65
5. The Incoherence of Meaning Irrealism 69
6. Wittgenstein's own Views on Rule-Following 76
7. Objectivity Secured 85
8. The Charge of Quietism 92
9. Is There Anything Left to Do? 96
10. Conclusions 99

CHAPTER 4. A FRAMEWORK FOR ETHICAL OBJECTIVISM 101

1. The Limited Lessons So Far 101
2. Wittgensteinian Realism and the Anti-Objectivist Challenge: 'Queerness' 103
3. Wittgensteinian Realism and the Anti-Objectivist Challenge: Convergence in Judgment 106
4. 'Response-Dependence' and Objectivity 107
5. What Role for Width of Role? 113
6. Conclusions 121

CHAPTER 5. AGAINST ONE VIRTUE-CENTRED THEORY OF ETHICAL DELIBERATION 124

1. One Model of Ethical Realism 124
2. Uncodifiability and the Limits of Internal Criticism 127
3. Codification and Deliberative Rules 135
4. Normativity and Substantive Disagreement 148
5. McDowell's Account of Practical Deliberation 157
6. 'Silencing' and the Virtuous Ideal 165
7. Conclusions 174

CHAPTER 6. TOWARDS A 'KANTIAN' ETHICAL COGNITIVISM 177

1. Taking Another Approach to Practical Rationality 177
2. Objectivity and Kant's Categorical Imperative 181
3. Towards Ethical Cognitivism 199
4. Delivering Ethical Objectivity 207
5. The Payoff 215
6. Some Concluding Remarks 228

Bibliography 235

LIST OF ABBREVIATIONS

Bibliographical references in this thesis will appear in the footnotes in the form of the author's name and the date of publication of the work. Full references can be found in the bibliography at the end of the thesis.

There is one group of references that forms an exception to this rule. In Chapter 3, frequent references will be made to the works of Ludwig Wittgenstein. In order to save on footnotes, references to these works will appear within the text and, in the interests of brevity, I shall use the following abbreviations to refer to them:

BB *The Blue and Brown Books* - Wittgenstein (1975).

CV *Culture and Value* (Second Edition) - Wittgenstein (1978a).

LA *Lectures and Conversations on Aesthetics, Psychology and Religious Belief* - Wittgenstein (1966).

LE "A Lecture on Ethics" - Wittgenstein (1965).

OC *On Certainty* - Wittgenstein (1977).

PI *Philosophical Investigations* - Wittgenstein (1976).

RFGB "Remarks on Fraser's Golden Bough" - Wittgenstein (1993).

RFM *Remarks on the Foundations of Mathematics* (Third Edition) - Wittgenstein (1978b).

Z *Zettel* - Wittgenstein (1967).

Full bibliographical details of these works can be found in the bibliography.

ACKNOWLEDGEMENTS

There are a number of people who have helped me in writing this thesis in a variety of ways. I am very grateful to Nicola Chalton, Peter Cornwell, Max de Gaynesford, Garrett Gullity, James Harris, Martin Kretschmer, Gerald Lang, Halvard Lillehammer, Shepley Orr, Bill Pollard, Peter Railton, Mark Sainsbury, Galen Strawson and Chon Tejedor-Palau for helpful discussion of, or comments upon, material contained in this thesis or its ancestor, my University of London M.Phil. thesis. Audiences at the Ockham Society, Oxford, the Oxford Philosophy Sub-Faculty Doctoral Thesis Seminar, and the 1999 British Society for Ethical Theory Conference, Durham, offered very helpful comments in response to presentations of earlier versions of various parts of this thesis. I owe a particularly great debt to those who have supervised my work over the last few years: Malcom Budd for my M.Phil. thesis, and Bill Child, Ralph Walker and, especially, James Griffin, for this thesis. Bill Child is owed a particular of debt of gratitude (maybe 'gratitude' is the wrong word!), as it was hearing a paper of his at the Chequer's Philosophical Society as an undergraduate in 1993 that first provoked me to start thinking about whether Wittgenstein's considerations on rule-following (views towards which I was sympathetic) really do drive one towards some kind of virtue-centred ethics (a position towards which I was, and remain, rather less sympathetic). I also gratefully acknowledge the receipt of a studentship from the Arts and Humanities Research Board of the British Academy for the years 1997-9. Work on this thesis would long since have foundered without the support of my parents, Peter and Hilary Cornwell, and of my partner, Clare Hawley.

Chapter 1

INTRODUCTION

1. The Challenge

At first blush, there might not seem to be much question that ethical discourse involves the stating of facts and that judgments about these facts can be objectively correct or incorrect.[1] Just as with uncontentiously factual subjects, we advance arguments for and against the morality of various courses of action. We also assess such arguments for their cogency and validity. The statements that we make about the ethical properties of situations, such as "John's lying is evil", seem to have the same assertoric form as statements that are about uncontroversially factual matters, such as the statement "The cat is on the mat".[2]

However, appearances can be deceptive, and these first impressions that ethics admits of objectively correct and factual answers have been forcefully challenged. Indeed, since the 1930s, the philosophical orthodoxy has been that this objectivist picture of ethics is deeply mistaken. The orthodoxy has been that, in fact, our ethical judgments do not involve the acquisition of beliefs about the ethical properties of situations, but instead

---

[1]There is a knotty tangle of issues about what is at stake here. Some possibilities are: whether ethical judgments can be objectively correct, whether ethical properties can be real, whether ethical judgments can be true or false, and whether ethical judgments can amount to knowledge of ethical facts. I hope that I can be forgiven a relatively loose initial presentation of the issue. In Chapter 2, I set about the task of getting a precise idea of what ethical objectivists/realists/cognitivists are disputing with ethical anti-objectivists/irrealists/non-cognitivists.

[2]Throughout the thesis I use the words 'moral' and 'ethical' more or less interchangeably.

amount to the expression of our attitudes or sentiments.[3] This denial of objectivity to ethical judgments has taken a number of forms. Some have claimed that our ethical thinking does not, in fact, rest on any objectivist presuppositions and that an anti-objectivist theory best explains ethical practice, as it actually is.[4] Others have argued that ordinary ethical thinking does, indeed, rest on such assumptions, but that in doing so it is radically in error, and thus the recognition of this error will entail a revision of our practices of ethical thought.[5] In the last decade, a third way of formulating ethical anti-objectivism has suggested that the surface appearance of truth-aptness for ethical statements is indeed correct. However, it holds that in ethics the type of truth predicate that is in play is much weaker than in, say, the physical sciences, and so ethical judgments cannot be objective.[6] What all of these forms of ethical anti-objectivism share is the view that there cannot be ethical properties that match up to the same standards of objectivity as those properties that feature in uncontroversially objective discourses, such as that of physical science. Hence, ethical judgments about such properties cannot be objectively correct in the way that our empirical judgments can.

This ethical anti-objectivism has been the orthodoxy, but over the last thirty years or so this orthodoxy has come under increasing threat from objectivist conceptions of ethics. These challenges have, of course, taken a wide variety of forms, but they can be

---

[3]See Darwall, Gibbard and Railton (1992), for a detailed account of how this dialectic developed over the last century. The view that ethical judgments express attitudes has, of course, an ancient lineage and finds its classic expression in David Hume's writings (see Hume (1975) and (1978)).

[4]Important presentations of this view are Ayer (1971), Ch.6, Hare (1952), (1963) and (1981), and Blackburn (1971) and (1984b), Ch.5 and 6.

[5]For this view, see Mackie (1977), Ch.1, and Harman (1977), Ch.1.

[6]See Wright (1996) for this approach. Wright (1992), *passim*, develops this as a general approach to taxonomizing disputes about realism and anti-realism.

seen as falling into two very broad categories. The first challenges the idea that ethical properties cannot pass muster alongside the uncontroversially naturalistic properties of science, and argues that, since they can, the reality or objectivity of ethical properties, and judgments about them, is not impugned. The other broad strategy is to challenge the idea that the scientifically-inspired conception of objectivity, which grounds ethical anti-objectivism, is in any way compulsory (or, more ambitiously, even coherent). My interest in this thesis is with objectivist responses of this second kind.

Again, there are a variety of ways of pursuing this very broad strategy, but the two that I shall be most concerned with in this thesis find a large part of their inspiration in two of the great figures of philosophical history: Aristotle and Immanuel Kant. The first of these strategies finds in Aristotle a hostility to the codification of rationality, which challenges the dominance of the scientific conception of objectivity. It also finds in Aristotle the basis for a positive account of ethics that combines objectivism with a recognition of the role of human nature and sentiment. This approach has been reinforced, in the work of several writers, by an appeal to Wittgenstein's views on the objectivity of meaning, which also challenge the dominance of a scientistic conception of objectivity.[7] The second strategy adopts Kant's appeal to thin norms of rationality as a means of vindicating the objectivity of judgments, and it maintains that this is a conception of

---

[7]See McDowell (1979) and (1981a), and also Lovibond (1983), for clear examples of this approach.

objectivity that can be fruitfully applied to ethics.[8] It claims that this is an adequate alternative conceptualization of objectivity to that offered by the scientific picture.

These two approaches - the virtue-centred Aristotelian one and the Kantian - have traditionally been seen as incompatible and antagonistic. This is because the first stresses the uncodifiablity of rationality and the central role of human nature in reason, while the second rests rationality on obedience to rational norms and a transcendental rational subject detached from the empirical world.

2. A Preview of the Project

This thesis has three main aims. The first is to argue that we can, indeed, find in Wittgenstein's discussion of rule-following compelling reasons for rejecting the scientifically-inspired conception of rationality and objectivity that seems to motivate the various forms of ethical anti-objectivism. Such a conception, it will be argued, is of very dubious coherence, and therefore the failure of ethics to match up to it can hardly constitute a problem for ethical objectivity. But this still leaves us with the positive task of producing a coherent account of what makes an ethical judgment objectively correct. So the second goal of the thesis is to argue that the virtue-centred approach developed from this starting point is not adequate to the task of accounting for ethical objectivity. There is no compelling reason to move from the Wittgensteinian considerations about rule-following and objectivity to a virtue-centred approach to first-order ethics. Worse

---

[8]This approach takes various forms. Some try to stay fairly close to Kant's original thought. Examples of this are Herman (1993), O'Neill (1989), Nell (1975), and Korsgaard (1996a) and (1996b). Others have taken an approach that is a clear descendant of Kant's, but is at some distance from it, for example 'contractualists' such as Rawls (1971), (1980) and (1985), Nagel (1991), and Scanlon (1982).

still, I shall argue, this approach allows itself too modest resources to provide a conception of objectivity that has any real bite in discriminating correct and incorrect ethical judgments. The third and final goal is to argue for a development of the Kantian approach to objectivity. This is not only compatible with a broadly Wittgensteinian conception of objectivity, but, when placed within such a framework, is able to overcome some of the problems that have traditionally dogged Kantian ethics. The result, I shall suggest, is an ethical objectivism that is successful in explaining what it would be for ethical judgments to be objective and that offers a plausible model of practical deliberation.

The road ahead will take the following course. In the next chapter, Chapter 2, I shall try to clarify what is at stake in the debate over the realism or objectivity of ethics. We will try to extract from the various challenges to ethical objectivism some idea of what the key problems are that anti-objectivists think undermine objectivism. I shall argue that the core of the anti-objectivist position is a scientistically-inspired conception of rationality and objectivity. According to this conception, judgments are sanctioned as objectively warranted from a perspective radically external to the conceptual practices of those making the judgments. Chapter 3 consists of an explication and examination of Wittgenstein's discussion of following a rule, in which such a conception comes under sustained and, I shall argue, fatal attack. I shall contend that Wittgenstein does not undermine the possibility of a reasonable standard of objectivity. Wittgenstein's grounding of objectivity in the practices of rational agents makes it natural to look for a form of ethical objectivism that is 'response-dependent'. On such an account, the extensions of ethical concepts are regarded as partially determined by the responses of rational ethical agents. Whether or not such a theory can deliver true objectivity will depend crucially on

the theory of rationality that determines which responses can appropriately determine the extension of ethical concepts.

In Chapter 4, I shall lay out the bare bones of this way of conceptualizing response-dependent forms of objectivism. The argument so far has developed from a discussion of Wittgenstein's views on objectivity. Thus it is logical to consider John McDowell's suggestion that the rule-following considerations imply that the only viable way of filling out the theory of rationality at the heart of the response-dependent objectivist model is along virtue-centred lines. I consider this suggestion in Chapter 5 and argue that Wittgenstein's rule-following considerations do not limit the options as severely as McDowell suggests. More damagingly for this form of neo-Aristotelian, virtue-centred ethics, I argue that its resources are too modest to provide anything that genuinely amounts to objectivity and that it suggests a deeply implausible ideal of practical deliberation.

The task that now lies before us is to work out an alternative, and more successful, way of developing the response-dependent objectivist account. In Chapter 6, I argue that, contrary to what many would suppose, the Kantian approach might offer a chance of meeting this challenge. It is an attractive way of carrying on the project because it shares with the Wittgensteinian conception of objectivity a rejection of the codifiability of rationality and a rejection of a reductionist or foundationalist grounding of rationality. It also holds out the promise of delivering a strong form of ethical objectivism, but does so by appealing to very modest additional conceptual resources. The Wittgensteinian framework turns out to provide a favourable environment in which the Kantian approach can be developed to overcome some of its traditional problems and deliver genuine

objectivity for ethical judgments. I also briefly draw out some of the further advantages of this approach as a model of practical deliberation.

3. A Few Methodological Remarks

Before getting properly started, it would be helpful to clarify a few methodological points about how the project is to be executed. Given the heady mix of famous dead philosophers (one of whom (Wittgenstein) had a rather unorthodox approach to philosophical method) and modern analytical meta-ethics it would be easy to lose sight of the rules of the game.

(a) Morality and moral theories

Attempts to connect together meta-ethical accounts of the grounding of ethical judgment with accounts of first-order ethical deliberation can take the form of moral theories. Such theories hold out the prospect of a route from metaphysical truths right through to guidance on how to act. For this reason, such theories are hopelessly overambitious and have been the subject of much warranted criticism.[9] This thesis does *not* try to present such a theory, and it should not be read as doing so.

Instead, the thesis develops an account of key features of ethical practice with a much more modest and realistic goal in sight. Given a highly complex area of thought (and ethics is undoubtedly such), developing a theory as a kind of working model can be one of the more fruitful ways of venturing into terrain that would otherwise be hopelessly confusing. A simplifying model can render partially tractable otherwise intractable conceptual challenges. Because it is necessarily simplified, such a model will be inaccurate.

---

[9]See Williams (1985) for an influential attack on the over-ambition of moral theories.

But as long as we are aware of what we are doing, are aware of this simplification and are prepared to revise, or even abandon, our model, this approach does not seem to be desperately unhealthy. Indeed, the approach seems further to recommend itself if we (rightly) reject the idea that we are trying to develop a 'first philosophy' that is to provide a firm foundation from which the rest of our conceptual scheme can be built. Rather, the task of philosophy seems to be to trace the links between our concepts. This task will issue in conclusions that are provisional and conditional in nature, telling us what concepts we can, and cannot, hold onto, given our starting points and our understanding of other concepts. On this way of thinking, these very starting points will be open to revision, or even abandonment. Given this framework, we can think of a theory as a tool for tracing out the links between certain concepts and for exploring the ways in which these concepts can be revised and the implications of any such revisions.

So, my aim is *not* to develop a 'first philosophy' for ethics. Rather, the aim is to use the Wittgensteinian conception of objectivity and a Kantian approach to ethical deliberation as a framework through which to explore the relations between the concepts of, for example, objectivity, the uncodifiability of rationality, truth in ethics and the categorical nature of moral requirements. The success of this project will be measured by its internal coherence and by how fruitful it is in shedding new light on these conceptual connections.

It seems to be particularly important to issue this disclaimer because of the nature of the two philosophical influences that will figure prominently in my argument. Kant is often seen as being a builder of theories, so I want to warn against the idea that my deploying Kantian ideas means that I, too, am engaged in such a task. A suspicion that this was my goal would provoke those who are, with considerable justification, hostile to

moral theories. On the other hand, Wittgenstein is an arch-opponent of philosophical theory-building, and my co-opting of some key ideas from his work would provoke immediate suspicion of a deep incoherence in my project if I was thought to be engaged in such theory-building.

(b) Is the project arbitrary?

It might be suggested that the line of argument that I have outlined has a rather arbitrary and quixotic character, in particular with respect to the positive account of ethical objectivity that I offer - is it not just a random and unmotivated cobbling together of disparate views simply for effect? In response to this charge, I would say that meta-ethics is an increasingly complex area with a myriad of possible approaches. To have systematically worked through all the extant options, arguing against them one by one, until my approach alone remained standing, would have been an impossibly labourious task leaving me no space in which to explore my positive proposals.

The argument against non-cognitivism, rooted in the Wittgensteinian conception of objectivity, is, I believe, a hugely powerful one that strikes at the heart of most forms of ethical non-cognitivism and anti-objectivism. It will probably also undermine those ethical objectivists who accept the same basic picture of reality. However, beyond this, if my arguments against the highly restrictive form of neo-Aristotelianism that McDowell develops are correct, then a large area of conceptual space could be left open for alternative ways of developing an objectivist position. I do not pretend that the option that I choose to pursue is the only one, nor do I provide anything but the most cursory of negative arguments as to why these other options are not to be favoured - it is open to others to develop alternatives within the conceptual constraints that have been unearthed.

However, I do think that the option that I have chosen promises to deliver a lot by way of objectivity, but requires us to move a surprisingly short distance away from our Wittgensteinian starting point. If nothing else, conceptual economy speaks in its favour. Given my remarks above about the role of theories as tools to help us elucidate the relations between complex and ill-understood concepts, I believe that something useful can be achieved, even without a systematic elimination of all the alternatives.

(c) Wittgenstein and 'Wittgensteinian', Kant and 'Kantian'

There is always a danger that in using the arguments of great philosophers one develops them to such a point that their original parentage comes in to question, and the coherence of their application may become doubtful.[10] I believe that I manage to avoid these dangers. My discussion of Wittgenstein's rule-following considerations involves a fair amount of exegesis, but my primary aim is to bring out the independent plausibility of the position advanced. Thus, whether or not the position is actually Wittgenstein's is not centrally important. What matters is that the argument (from wherever it has been unearthed) works. It is an argument that is clearly heavily inspired by material of Wittgenstein's, and thus should qualify as 'Wittgensteinian', even if it is not Wittgenstein's.[11]

With respect to Kant, I do not offer any exegesis but deploy some ideas that emerge from an interpretation of Kant's views. Again, I hope that I succeed in establishing the independent plausibility of these ideas and their coherence with what has gone before, so that it ultimately does not matter too much if they are not actually Kant's. My account

---

[10]See Engberg-Pedersen (1995), pp.122-4, for a useful discussion of how things can go wrong in this regard.

[11]But, again, it is not crucial if it does not.

deploys a version of Kant's notion of the Categorical Imperative and also seeks, as Kant did, to elucidate the connections between rationality, objectivity, and the categorical nature of moral requirements, and so it would seem to warrant the label 'Kantian'. But, again, if in the end it does not, this will not undermine the ultimate plausibility of the position. The plausibility of such a position is only undermined if the argument for it essentially rests upon some other part of the great dead philosopher's system that has either not been endorsed or has actively been rejected. I am confident that this is not the case with my appeal to either Kant's or Wittgenstein's ideas.

The other way in which it might be thought that the project could come to grief is if the further development of my argument, after the deployment of the Wittgensteinian or Kantian materials, conflicts with those materials. But really this point only reminds us of one's responsibility for the coherence of one's own argument, irrespective of from where the bits that make it up might have been borrowed. It might be thought that a problem also arises if the further development of the argument brings me into conflict with other actual views of Wittgenstein or Kant. In both cases this probably happens. Wittgenstein never explicitly linked his discussion of rule-following to ethics, but from the little that he did write on ethics it would be fair to say that the position that I ultimately develop is pretty unWittgensteinian.[12] But this is not a flaw in my position because I am not trying to develop an exegesis of Wittgenstein's ethical views, but seek to examine how another area of his thought might be developed to shed light on ethics. Similarly with Kant, my position will ultimately involve elements that are quite foreign to Kant's own position: an acceptance of the role of human nature in the grounding of rational practice;

---

[12]Wittgenstein's views on ethics can be gleaned from *CV, LA, LE* and *RFGB*. Johnson (1989) tries to develop a systematic account of ethics based on these materials.

no appeal to the transcendental self as the source of rational agency; and the acceptance of the existence of ethical properties and of ethical agents making judgments about them. Again, such conflicts 'downstream' (as it were) of the deployment of a great philosopher's ideas are not a problem, so long as one is not claiming to be doing pure exegesis, which I am not.

(d) Global and local anti-objectivism

The debate between ethical objectivists and anti-objectivists, as it appears in the literature of analytical moral philosophy, rests on the assumption that it is possible for judgments within some discourses to be objective. Ethical anti-objectivists do not deny that there are objective judgments; in fact, quite the opposite - they argue that judgments within a certain favoured discourse (for example, physical science) have just this status and the judgments in certain other discourses (ethics, aesthetics, politics, for example) lack such a status. For such theorists, it will count as a serious problem for their position if it is found that the anti-objectivism that they argue for with respect to their target discourse somehow globalizes to challenge the objectivity of judgments in *any* discourse.

There is another form of anti-objectivism or irrealism about ethics that would see the failure of ethical judgments to achieve objective correctness merely as a particular case of a global irrealism that denies that objectivity is possible in any discourse. Such a view is not one that I shall address directly, although a clear consequence of the argument in Chapter 3.5 will be that such a view is incoherent. A major motivation for such a view is a disenchantment with the scientistically-inspired conception of objectivity, but the advocate of such irrealism shares with the advocate of that conception the belief that it is the only proper conception of objectivity. This irrealist, however, draws the conclusion

that, since this conception is incoherent, there can be no objectivity. The argument that I find in Wittgenstein suggests that this inference is illegitimate.

(e) Locating the goalposts

It is sometimes a little hard to see what would constitute a successful argument for an ethical objectivist (or realist or cognitivist). Anti-objectivists have a relatively easy job; for them it can suffice to provide some entirely a priori argument as to why some feature of ethical judgment is incompatible with some key aspect of objectivity. Their opponent has a harder job. For them it does not suffice to provide an a priori refutation of the anti-objectivist's position, for that seems merely to open up the possibility that ethical judgments can be objectively correct - it does not show that any actually are. It would be a pretty limited form of ethical objectivism that was compatible with all our ethical judgments actually turning out to be incorrect. On the other hand, it can hardly be demanded that objectivists not only refute their opponents' arguments, but also show that all ethical judgments are either correct, or incorrect for some explicable reasons, and that a sufficiently large class of them are, in fact, correct. The first task would seem to be impossible to execute. It is also of dubious coherence, since it is quite compatible with objectivism that for some range of judgments it is indeterminate whether they are correct or not. We might simply be lacking in the conceptual resources to determine the answer. The second task seems entirely under-specified, since it is open to question what the appropriate proportion of correct judgments is - all, 99.9%, 90%, 80%, or what proportion?

I shall take it that objectivists will have done a reasonable job if they refute their opponents' arguments and offer a way of determining the correctness of ethical judgments

that can yield determinate results in a range of cases. They will also need to give some explanation, compatible with this, of those cases where success is not achieved.

So, with these methodological points made, we can now proceed with advancing the argument. The first thing to do is to develop a much more precise formulation of what is at issue between ethical objectivists and their opponents, and it is to this task that we move in the next chapter.

Chapter 2

ETHICAL OBJECTIVITY: WHAT IS AT STAKE?

1. Some Terminological Preliminaries

The question of what is at stake in disputes about ethical objectivity is notoriously obscure, and it is often muddled up with other, more or less cognate, issues - moral realism, cognitivism, the role of truth in ethics, and the motivating force of ethical judgments, for example. Any attempt to get a precise fix on the issue will necessarily ignore some currents that run through the debate and will not satisfy all parties. But the key issue can really be seen as that of what form (if any) justified ethical judgments take. An ethical objectivist will say that an ethical judgment is justified if it matches up to an objective standard of correctness. This does not tell us very much, so the next question will be what form this standard might take. An obvious candidate as an answer is that an ethical judgment is objectively correct if it is substantially true.[1] A corollary of this is that ethical judgments are beliefs and reaching an ethical judgment is a matter of forming a belief. If the belief is true, then one might say that the reaching of such a judgment even constitutes the acquisition of knowledge, provided, of course, that whatever further conditions required by one's pet theory of knowledge are also satisfied by the belief. I shall label this view that ethical judgments can be true or false ethical or moral 'cognitivism'.

---

[1]This is still pretty vague, but I mean by this that it is true in a more than merely minimalist sense or in the sense in which a quasi-realist might accept that ethical judgments can be true. What this means will become clearer shortly.

Some would see this position as the core of ethical objectivism, but I want to use that term in a rather wider sense that includes any theory that sees ethical judgments as capable of being objectively correct or incorrect. This allows one to classify as objectivist a theory such as Kant's, which sees moral judgment as answerable to objective rational norms. Kant's theory is generally regarded as being strongly objectivist, whilst eschewing any form of ethical cognitivism, and it seems appropriate that our definition of ethical objectivity should be able to group Kant with the objectivists.[2] There certainly seems to be something intuitively rather strange about saying that Kant's theory is anti-objectivist or non-objectivist, so I shall take objectivism as covering Kant's theory.

What I have labelled ethical or moral 'cognitivism' is often referred to as 'moral realism'. I shall try to avoid doing this for two broad reasons. The first is that, as David Wiggins has pointed out, the term 'moral realism' can provoke some confusing and unwanted associations. It can become associated with moral absolutism, with scientific realism, with debates between realism and idealism/mentalism in metaphysics or between platonic realism and conceptualism in mathematics, and, finally, with Dummettian realism versus anti-realism debates.[3] The second reason is that some of those who are most ready to apply the label 'moral realist' to their views are willing to accept a strong form of moral cognitivism. This goes along with a certain theory of ethical deliberation that sees ethical judgment as a process of perceiving moral facts, or, less controversially and more charitably, perceiving the moral features of situations.[4] As the theory that I develop stops

---

[2]In this I depart from the otherwise very useful definition of 'objective' in Wiggins (1996a), p.35.

[3]Wiggins (1990-1), pp.62-3.

[4]See, for example, Dancy (1993), McDowell (1979), (1981a), (1985b) and (1995a), McNaughton (1988), and Platts (1979), Ch.X.

short of such a full-blooded cognitivism, I shall steer clear of the term 'moral realism', lest others, like me, find that it has these associations. Ethical cognitivism, as I shall use the term, need not be seen as anything other than neutral about the form that the process of ethical judgment can take.

While we are in the process of clarifying terminology, it is worth noting that I have made no mention of subjectivism. The reason for this is that I do not want to regard subjectivism as the rejection of objectivism (positively defined), nor, vice versa, objectivism as the denial of subjectivism (positively defined). Rather, I take it that the subjective is to be defined positively, and independently of the objective or objectivism, as (roughly) pertaining to the states or responses of subjects.[5] It will be a matter for investigation, not definition, whether or not a subject matter can be objective *and* subjective. In this chapter, I am mostly concerned with issues concerning objectivity. The account of objectivity that I develop in the following chapters will be one that might be seen as subjectivist in a very weak sense, but, if it succeeds, will nonetheless be a form of relatively strong objectivism.

There is, in the literature, often a failure to distinguish between the prospects of objectivism as it pertains to axiology - that is, the theory of evaluations - and as it pertains to the theory of practical reasons.[6] It is quite possible that the prospects for objectivism will be very different in these two different areas, and there may even be room for variations within, as well as between, them. This should not unduly worry those whose interest in examining the ethical domain stems from a genuine concern for a better

---

[5]Again, see Wiggins (1996a), pp.36ff, for this fruitful way of defining the term 'subjective'.

[6]See Wiggins (1998), pp.95-6. Crispin Wright, in Wright (1988b) and (1996), for example, seems insufficiently aware of this distinction.

understanding of its resources. It might, however, disturb those who are more interested in squeezing ethics *en masse* into the framework of a generic realist versus anti-realist debate. My interest will be in developing an objectivism that applies primarily to the theory of practical reasons.

2. Routes through the Maze

With our terminology in better order, we can get the debate moving. Starting with our very rough definition of objectivism, I shall attempt, in this chapter, to circle in on a progressively tighter account of the issues for objectivism. In this way, we will develop a clearer picture of what counts as an objectivist theory. This process will issue in conclusions that are, to an extent, provisional. The process of further circling in on our target will continue over the next two chapters as I develop a positive account of objectivity, first in general and then as it might obtain in the ethical domain.

Given the dominant position of non-cognitivism or anti-objectivism in the debate, I shall try to take on my opponents in their own territory. I shall try to tease out what notions of objectivity ethical non-cognitivists hold and why they think that ethical judgments cannot be regarded as objectively correct. There are various ways in which ethical non-cognitivists have tried to set up their challenge to ethical objectivism, and I shall try to examine some of these broad strategies, without claiming to offer a detailed discussion of individual writers. In the following two sections I shall merely lay out these approaches to ethics without examining their reasons for rejecting objectivity in this domain. Section 5 will then try to draw out the aspects of ethical objectivism that they see as objectionable. The remainder of the chapter will be concerned with sifting through these challenges to try to arrive at a clearer idea of which parts of the anti-objectivist challenge

involve legitimate demands and which might rest on more questionable assumptions about the nature of objectivity. I shall argue that the non-cognitivist's conception of objectivity, which makes the supposed objectivity of ethical judgments so problematic, takes a certain understanding of physical science as its paradigm.[7] On this view, judgments concerning a domain achieve objectivity only in so far as they are accessible from a view from nowhere, or concern the 'fabric of the world' or primary qualities. The metaphors are powerful, but, I shall argue in the next chapter, nothing coherent lies behind them, because the picture of reality and objectivity that they are supposed to illuminate is ultimately of very dubious coherence.

Although all the anti-objectivist formulations of the debate that I shall discuss give a central role to the notion of truth, their importance is not simply restricted to ethical cognitivism and realism. What they find problematic in ethical realism or cognitivism are largely features that these positions share with wider forms of ethical objectivism, and hence they suggest challenges to all forms of objectivism.

Much of the force of non-cognitivism lies in its powerful and intuitively compelling picture of objectivity, so, once this is challenged, the position loses much of its appeal. Objectivists still face the task of clarifying their own account of objectivity and satisfactorily showing that ethical judgments do, to at least some extent, match up to it. This is a task that will take up the latter part of the thesis. Of course, this enterprise might fail, and, even shorn of their master argument, non-cognitivists might, in their positive accounts of how ethical judgments exhibit the discipline that they seem to, have the most

---

[7]I say 'a certain understanding of physical science', because the point is not to question the objectivity of science, but to challenge a certain account of what that objectivity consists in and to claim that this understanding is, in fact, a *mis*understanding.

compelling account of the seeming objectivity of ethics. But until objectivists do fail in their task, non-cognitivists will seem to be attempting to provide an elaborate answer to a question that rests on a misunderstanding.

3. Equating Objectivity with Truth-Aptness

An obvious candidate as an answer to the question, What makes a discourse open to objective assessment?, is that the sentences in it are apt to be assessed as to their truth or falsity. The connection between truth and objectivity will depend on what conception of truth is in operation and on what role truth-aptness is to play in the dialectic between objectivist and anti-objectivist, or between realist and anti-realist.

Traditionally, the connection between truth and objectivity has been regarded as relatively straightforward: if a discourse is genuinely truth-apt and at least some of its sentences are in fact true, then the discourse is objective; if not, then it is not objective. Thus, broadly speaking, the anti-objectivist has tried to show one of two things about the discourse in question. The first alternative sees the ethical anti-objectivist trying to show that the discourse does, indeed, proceed as if its sentences are assessable as true or false, when, in fact, the sentences are universally false. The classic example of this approach is J.L. Mackie's treatment of ethics. His arguments from 'queerness' and relativity are used supposedly to show that there could be no such things as the moral facts or properties that ethical judgments purport to reflect, and hence that all ethical judgments are false because they fail to refer.[8] This approach takes seriously the apparently objective and truth-stating surface appearance of sentences expressing ethical judgments.

---

[8]Mackie (1977), pp.36ff.

Alternatively, it can be argued that, although a certain discourse appears to be assertoric in nature when its surface features are examined, in fact, deep down, its sentences have some sort of non-assertoric function. Here the classic examples are R.M. Hare's treatment of moral judgments as universal prescriptions and Simon Blackburn's quasi-realist treatment of ethical, probabilistic and modal discourse.[9] According to the latter, these sorts of judgment are not really responses to features of the world, but are, rather, the result of a process of Humean "gilding and staining" where we project our attitudes onto the world.

In response to an error theory, such as Mackie's, two lines of attack lie open to objectivists. First, they can attack the tenability of the conception of reality that gives rise to the error theorist's charge that the discourse is systematically false.[10] Or, secondly, they can raise a dilemma for the error theorist's account of how we can proceed after we have discovered that the discourse is in error. On the first horn of the dilemma, the error theorist takes the conclusion that a discourse is systematically false as a sign that, in all good faith, we ought to cease using it. This is obviously a very radical option, particularly given the centrality to our lives of the discourses that have attracted error theoretic accounts, for example, ethics and mathematics. The radical nature of this approach makes one want to be doubly sure that the error theorist's conception of objectivity really is compulsory. Unsurprisingly, error theorists have been reluctant to take this approach. This leads us (or rather, them) to the second horn of the dilemma. Here, they try to distinguish judgments that are appropriate and inappropriate on the basis of some norm other than truth. But now the problem is that, if such a norm can be applied, then why can

---

[9]See Hare (1952), and Blackburn (1984b), Ch.6, and (1993), especially Essays 3, 8 and 10.

[10]This is John McDowell's strategy in McDowell (1983) and (1985b).

we not re-construe our conception of truth in terms of this norm, rather than stick with the initial conception of truth that gave rise to the error theorist's conclusion?[11] This option turns the error theorist into something very close to a prescriptivist, expressivist or quasi-realist non-cognitivist.[12]

In contrast to the error theorist, the expressivist accepts that the surface syntax of sentences in the target discourse is assertoric, but insists that deep down something different, and non-assertoric, is going on. Because they accept the burden of maintaining appearances, expressivists or projectivists must explain how non-assertoric machinery can generate the discipline that is typical of apparently truth-apt discourses (like modal, ethical or probabilistic ones). This discipline seems to be easily explained if these discourses are taken to be genuinely assertoric. The emotivist or projectivist may have a seemingly plausible account of simple ethical judgments, such as "Tom's lie was wrong", according to which this is seen, not as the statement of a fact, but as the projection of an attitude of displeasure towards Tom's lie. However, moral language is marked by a host of more complex sentences involving logical constants, which seem to operate just as successfully as similar sentences do in uncontroversially factual discourses. Thus, to use a rather hackneyed example, the following *modus ponens* argument seems to be valid:

If stealing is wrong, then encouraging others to steal is wrong.

Stealing is wrong.

Therefore, encouraging others to steal is wrong.

---

[11]The dilemma is posed by Wright (1996), pp.2-3.

[12]Or, if the norms that they highlight are sufficiently robust, they might even end up with something close to a form of objectivism or cognitivism.

We have a satisfactory explanation of the validity of *modus ponens* where the sentences express truths. But what is the justification of *modus ponens* and other forms of logical argument when the function of sentences is, as non-cognitivists claim is the case with ethics, to express attitudes? At this point in the debate it suffices to say that an uncontroversially convincing explanation has not yet been offered.[13]

Clearly, expressivism will collapse into error theory if these technical problems cannot be resolved. But such technical niceties will seem somewhat beside the point if non-cognitivism rests on a fundamentally flawed metaphysics, and if ethical objectivism can form a coherent alternative. It is open to the objectivist, just as it was in response to the error theorist, to challenge the conception of objectivity that the non-cognitivist holds, and this will be the strategy that I shall pursue against the various varieties of ethical anti-objectivism.

At the end of this whistle-stop mapping of the argumentative positions, it can be seen that the argumentative burdens that remain, when the debate is construed in this way, are as follows. Error theorists have the job of showing how to stop the norms that they still think we can use in ethical thought from being reconstructed into a conception of truth. Expressivists have the job of giving a plausible, and independently warranted,

---

[13]The *locus classicus* for this problem is Geach (1965). Blackburn offers an expressivist answer to the so-called Frege-Geach problem in Blackburn (1984b), Ch.6. S.L. Hurley argues (Hurley (1984)) that the non-cognitivist pursuing Blackburn's line is forced to abandon the (independently plausible and attractive) idea that there should be a unified treatment of cases of *modus ponens*. For, while it is normally taken that a failure of *modus ponens* reasoning is a logical, purely rational failure, Blackburn suggests that it is a moral failure, that is, having an inappropriate complex attitude to situation. This radical move seems to be in need of independent support if the expressivist project is to convince. Also required are: (a) a satisfactory account of what an *inappropriate* complex attitude is, and (b) some account of what connects the complex attitudes in various cases of *modus ponens* reasoning, so that they can be seen as examples of one unified phenomenon, rather than a set of unrelated attitudes to particular instances of reasoning. See, also, Crispin Wright's discussion of the problem in Wright (1988b), pp.31-4.

account of the supposed non-assertoric nature of the underlying mechanics of ethical discourse. Objectivists have the job of showing that error theorists and expressivists have an illegitimately restricted view of what counts as objective. My concern will be with this last task, since success here will undermine the motivation for pursuing the tasks that have been allotted to error theorists and expressivists. Their failure to complete these tasks would not be an argument for ethical objectivism, since it would still be possible to embrace ethical nihilism if objectivists also fail in their task.

4. Minimalist Truth and Objectivity

As we saw in the last section, setting up the debate in terms of truth-aptness versus non-truth-aptness left ethical anti-objectivists, of whatever stripe, with a fairly demanding positive task to fulfil. Anti-objectivists might, therefore, be tempted to find a way of setting up the debate under which they do not have to adopt this sort of burden and objectivists are left doing all the work. There is an alternative way of construing the relationship between truth and objectivity, so that the issue is not one of truth-aptness versus non-truth-aptness, but one of what *kind* of truth predicate can be applied to a particular discourse. This allows anti-objectivists or anti-realists an easy way of respecting the surface assertoric features of, for example, ethical discourse, because they can accept that the discourse is, indeed, truth-apt, but not in any realism- or objectivism-supporting manner. In this way, anti-objectivists might hope to get their case going without taking on anything like the expressivist's obligation to give a non-assertoric reconstrual of the apparently assertoric features of ethical discourse.

This way of setting up the debate has been suggested by Crispin Wright.[14] Wright puts forward a conception of truth - minimalism - that requires for truth-aptness only that a sentence possess assertoric content. The notion of assertoric content is glossed in terms of a sentence's meeting certain syntactic and disciplinary constraints, so that:

> [E]ssentially, sentences are assertoric which are capable of significant embedding within constructions such as negation, the conditional, and in contexts of propositional attitude, and whose use is subject to acknowledged standards of warrant.[15]

The existence of a truth predicate defined for sentences meeting this standard will, claims Wright, mean that such a predicate would be governed by the Disquotational Scheme:

> 'p' is true if and only if p.

Contra a deflationist conception of truth, this entails that:

> [T]he word 'true' will record a norm governing assertion and belief-formation which is distinct from assertibility, i.e. warrant by whatever standards inform the discourse in question, and that its compliance or non-compliance with *this* norm can hardly fail to be reckoned to be a substantial property of a statement.[16]

Wright thinks that we can pick out a number of principles or 'platitudes' that explain the meaning of the word 'true' in a uniform way across different areas of discourse. In addition to the Disquotational Scheme, he suggests the following platitudes: the thesis that to assert is to present as true; that to every truth-apt content corresponds a truth-apt negation; that a content is true just where it corresponds to the facts, depicts things as they are, and so on; that truth and justification are distinct; that truth is absolute; and that truth is stable.[17]

---

[14]See Wright (1988b), (1992) and (1996). Because of his concern with general realist versus anti-realist debates, Wright sees his opponent as a moral *realist*. But for our purposes, Wright's attack on realism can be seen as essentially an attack on ethical objectivism.

[15]Wright (1996), p.5.

[16]Wright (1996), p.6.

[17]Wright (1996), pp.6-7.

On this way of construing the debate, the next question is, What would a realist have to say about a discourse that is not already covered by the minimalist conception of truth and its associated platitudes? To sharpen this question a little, we need to introduce one last bit of Wright's terminology, namely the notion of 'superassertibility'. A statement is superassertible "if it is assertable in some state of information and then remains so no matter how that state of information is enlarged upon or improved".[18] Superassertibility is constructed out of assertibility and supports all of Wright's platitudes listed above. So, the question becomes one of why we should not identify truth in a certain discourse with mere superassertibility.

Wright conceives of the moral realist as claiming that truth in ethics goes beyond superassertibility - that:

> [R]esponsibly to practice in [the relevant region of discourse] is to enter into a kind of representational mode of cognitive function, comparable in relevant respects to, say, taking a photograph or making a wax impression of a key .... [W]e put ourselves at the mercy, so to speak, of the standards of appraisal appropriate to the discourse in question (compare taking the snapshot or impressing the key on the wax).[19]

But Wright sees this conception as unsustainable in the case of ethical discourse for two reasons. First, he claims that it precludes the possibility of essentially contestable ethical concepts. Objectivism suggests that every moral dispute would be cast in terms that mean that one party is right and the other wrong, but Wright takes it that "a substantial body of the principles that inform our ordinary moral thought are essentially contestable".[20] To put the point differently, the realist would have to claim that ethical judgments exhibit, what Wright dubs, Cognitive Command, but the existence of essentially contestable ethical

---

[18]Wright (1996), p.10.

[19]Wright (1996), p.13.

[20]Wright (1996), p.15.

concepts renders this claim implausible.[21] Secondly, the ethical realist cannot make good the notion of correspondence to facts that are robust and independent of our practices in the way that a photographed scene is independent of the photograph.

As I said earlier, the purpose of this chapter is not to come to any hard and fast arbitration of whether any form of objectivism is true, but to get the question right. So in this context we must ask, Are Wright's challenges to the moral realist or objectivist in good order? In response to this I shall say something about Wright's first objection in Section 6 and about his second objection in Section 7, but let me first make a quick point about the latter.

Wright thinks that, for realism to be a distinct position, realists must have some robust notion of correspondence that goes beyond the Correspondence Platitude that is affirmed by minimalism (that is, that a content is true just where it corresponds to the facts, depicts things as they are, and so on). This platitude is merely derived from the Disquotational Scheme and hence has "much less metaphysical commitment than realism supposes".[22] Wright backs up this claim with talk of photographs and impressions in wax. This objection to ethical realism seems off target in the context of the dialectic that Wright is trying to recharacterize. At best it is an objection to the tenability of realism *tout court*, not to specifically *moral* realism, since it amounts to an attack on the coherence of the idea that a substantive conception of correspondence could, in and of itself, yield a conception of truth, correctness, and so on. Worse than this, it is a conception of truth that many who take themselves to be realists reject. It is one that almost all the notable contemporary *moral* realists and cognitivists do reject, being influenced, as they often are,

---

[21]I shall have more to say on Cognitive Command below in Section 6.

[22]Wright (1996), p.12.

by Davidsonian views about truth and interpretation.[23] Such realists would accept the claim of Wright's minimalism that the idea of correspondence is to be understood simply as derived from that of truth.

Even if Wright's point about the moral realist's failure to render intelligible a robust form of correspondence or representation is unhelpful, there is a better point that he does make. What the notion of correspondence or representation brings out is the idea that the realist wants to offer some picture of moral judgment that is 'driven by the facts', where these facts are seen as independent of the practice of making those judgments, in the way that a scene being photographed is independent of the photograph. This seems a much better point, and one that might seem to lie at the very heart of intuitive ideas about realism and objectivity. We will return to it shortly in Section 7.

This way of formulating the debate has the advantage that, by respecting the truth-apt surface appearance of ethical discourse, it is able to avoid burdening itself with the task of revising moral language. All the work has to be done by the realist or objectivist to show that ethical discourse supports anything more than the minimalist notion of truth. Because this approach operates entirely at the level of philosophical logic and metaphysics, it cannot readily find a place for traditional Humean reasons for being an ethical anti-objectivist. One may, or may not, regard this as weakness of the approach.

---

[23]David Wiggins makes exactly this point when he writes, "Let me say at once that I am not going to call myself an anti-realist simply because I have long since rejected the idea that correspondence can be used to elucidate that of truth" (Wiggins (1996a), p.48). See also Wiggins (1980) and "Truth, and Truth as Predicated of Moral Judgements", in Wiggins (1998). Donald Davidson's rejection of correspondence conceptions of truth can be found, for example, in "True to the Facts" in Davidson (1984), and in Davidson (1990).

5. The Rejection of Objectivism

Irrespective of the way in which ethical anti-objectivists want to set up the debate, there seems to be some convergence in their views on the key problems facing ethical objectivism, and, I shall argue, there is a common source for their suspicions. These various objections can be grouped, very roughly, as follows:

(i) Ethical objectivism demands too much. This group of objections suggests that ethical objectivism would require an implausibly complex metaphysics because of the connection of ethical judgment to action. It is a feature of morality that moral judgments have a peculiarly close connection to reasons to act. Thus, if I judge that a course of action is right, then I have a reason (possibly even an overriding one) to pursue it. But the anti-objectivist argues that, if ethical judgment is a matter of acquiring beliefs, then this situation is highly problematic. For Mackie, this prescriptivity of moral judgment means that moral properties must have an intrinsic "to-be-pursuedness" that renders them hopelessly "queer", because they would have to be detectable by some special moral faculty of we know not what kind.[24] Similarly, Blackburn, and before him, Hume, argued that the function of belief is passively to reflect the world and not to change it, and thus a moral judgment cannot take the form of a belief without severing the link with motivation to act. For these philosophers, when agents make ethical judgments, they have a non-cognitive attitude, such as a desire, directed towards a situation. It is worth

[24]Mackie (1977), pp.38-42.

remembering that this sort of moral psychological objection to ethical objectivism does not obviously fit neatly into the purely metaphysical framework that Wright suggests.[25]

(ii) Ethical objectivity demands too little. This kind of objection can be usefully seen as coming in two (interconnected) varieties. The first - which we might dub the 'Proper Convergence Demand' - claims that, where there is a convergence in ethical judgment ethical objectivism cannot explain why this amounts to anything more objective than a mere following of the herd. Hence, it is not an objectivity-sustaining form of convergence. The second - the 'Explanatory Role Demand' - claims that ethical properties have no role in explaining how an agent comes to a moral judgment and, indeed, have no role in explaining anything else either. I shall explain these two demands in more detail and discuss how they are connected in Section 7.

(iii) Ethical objectivism is incompatible with the essential contestability of ethical concepts. This third, and final, group of objections develops the common intuition that there is simply too much disagreement about the application of ethical concepts for ethical judgments to count as objective. Mackie's 'Argument from relativity' is, in part, an argument of this form. Wright's concern that ethical judgments lack Cognitive Command is also broadly of this form. There is a close connection between this kind of objection to objectivism and those in the previous category. The mere fact of disagreement, on its own,

---

[25]However, the compatibility of expressivism with Wright's minimalist framework is argued for by Smith (1994b) and (1994c). Divers and Miller (1994) and Horwich (1994) reject Smith's view. Horwich (1990) shares Wright's view about the relationship between minimalism, the role of belief and ethical anti-realism.

does not undermine objectivism, unless there is some additional argument as to why this disagreement is, in principle, unresolvable.

This, then, is a very crude way of sketching the lines of attack that ethical anti-objectivists use against objectivism. It will become clear that there is a considerable degree of interconnection between these arguments. Indeed, this interconnectedness is, arguably, not an accident, for the best way of making these objections more concrete is to consider how they fail to apply to a paradigmatically objective discourse, namely that of physical science. According to a certain view of the situation, scientific judgments import no 'queerness' because they are not linked to any motivation to act. For the same reason, there is no problem in regarding scientific judgments as issuing in Humean beliefs. When subjects converge in their scientific judgment, this convergence is objectivity-supporting because the convergence is explained by something external to the belief system of the observers, namely properties of the material world. Connectedly, the properties of the physical world have a full causal explanatory role because, not only do they cause beliefs in observers, but they also interact causally with other properties. Finally, if there are disagreements about scientific judgments, then these are limited in number and can be attributed to inadequate data, vagueness in the concepts used or error on the part of one of the parties to the disagreement - all of which factors are potentially correctable.

Although the three groups of objections can be seen most clearly as objections to ethical cognitivism, they also raise problems for other, weaker, forms of ethical objectivism. If we take as an example Kant's theory, we can see that one who accepts the argument from queerness will find baffling the idea of reasons that are categorical and

entirely independent of a subject's inclinations or desires.[26] Similarly, Humean emotivists will find it deeply problematic that the finding of a contradiction in one's maxim to act could give rise to a desire. More fundamentally, they would also reject the idea that one could have a categorical reason to act that was independent of one's desires.[27] Crucially, as we will see later, it is not so clear that Kant's views face difficulties with the Proper Convergence Demand. But they seem likely to fare pretty badly in relation to the Explanatory Role Demand. Without satisfying the Explanatory Role Demand, it will appear that the Kantian position has few genuine resources to deal with the anti-objectivist concerns about relativism or lack of Cognitive Command. So, this very crude way of formulating the anti-objectivist's challenges does seem general enough to hit at the heart of forms of ethical objectivism broader than cognitivism.

So, how is one to deal with these challenges? In response to the first set of challenges concerning queerness and prescriptivity, objectivists could get themselves off the hook by denying any connection between ethical judgment and motivation to act.[28] Instead of taking this route, I shall argue that the assumption that reality, or our beliefs about it, are motivationally inert is a prejudice that goes along with a misguided conception of objectivity. The task of attacking that conception will be one for the next chapter, and the task of developing a positive account of how the prescriptivity of ethical judgments might be explained will be started in Chapter 5 and completed in Chapter 6.

---

[26]See Mackie (1977), pp.29-30.

[27]See the essay "Internal and External Reasons", in Williams (1981), for an influential and sophisticated neo-Humean defence of this claim.

[28]This would be to adopt 'externalism' about ethical reasons.

The idea that objectivism involves convergence in judgment of an appropriate sort - the Proper Convergence Demand - does seem to be an intuitively compelling one. However, I shall argue in Section 7 that the ethical anti-objectivist faces a dilemma in spelling out this idea. On the one hand, one can be fairly liberal in the resources that one allows a theorist to appeal to in filling out the idea of what counts as an appropriately reached convergence. On the other, one can restrict the grounding of this convergence to a responsiveness to properties that pull their weight in the world, as Wright's Width of Cosmological Role requirement does. But, then, I shall argue, the plausibility of this restriction depends crucially on the availability of the ethical anti-objectivist's favoured model of reality, and this will be something that comes under attack in the next chapter. If this motivation for the restriction is removed, then it is not clear why the ethical objectivist is not free to develop other ways of defining the criteria for appropriate convergence.

This leaves us with the worries about relativity and Cognitive Command. I shall argue in the next section that, for these to present a pressing problem for the ethical objectivist, requires a good a priori argument why disagreements are not, in principle, resolvable. The non-cognitivist's considerations about proper convergence might play this role, but, without such an argument, worries about relativity do not themselves add anything to the anti-objectivist's challenge.

6. Essential Contestability and Objectivity

The existence of ethical disagreement is often thought to pose a serious problem for ethical objectivism, but the precise nature of the challenge needs careful handling. Crispin Wright suggests that, in order for a discourse to sustain objectivity, it must exhibit

Cognitive Command. Unless attributable to vagueness in the disputed judgment or standards of acceptability or to differences in personal evidential thresholds, it is a priori that any differences of opinion must be due to cognitive shortcomings on the part of at least one of the parties to the dispute.[29] The conclusion that ethical discourse does not exhibit Cognitive Command is suggested by the fact that:

> [A] substantial body of the principles that inform our ordinary moral thought are essentially contestable, and that no rational or cognitive deficiency is needed to sustain the clashes on things like sexual morality, the value of individual freedom, the moral status of animals, and the ethics of suicide and mercy-killing.[30]

In order to evaluate the force of this popular claim, one needs to clarify a number of points, in particular how widespread such disagreement is and what its existence would prove.

The anti-objectivist claims that moral disagreement is widespread both within and between cultures. This tends to be taken as a self-evident truth, but if sheer quantity of disagreement is meant to challenge the objectivity of ethical judgments, the ethical objectivist is surely entitled to ask what the scoring system is. No one would deny that there are areas of intense moral dispute within and between cultures, but there are also huge areas of agreement within, and between, cultures. Anyone in a modern Western culture advocating slavery, paedophilia, absolute monarchy, or the killing of another human being without a very good reason (although there is considerable debate about what counts as a very good reason) would be the ethical equivalent of a flat earther. Similarly, it should be a startling fact for the anti-objectivist or relativist that a mode of moral thought that gives a central place to a role reversal test has found expression in a

---

[29]Wright (1992), p.144.

[30]Wright (1996), p.15.

huge variety of sources - the ethical thought of Confucius in China two and a half thousand years ago, the Ten Commandments, Jesus Christ's injunction to love your neighbour as yourself, the moral thought of the Hopi Indians, and, of course, contemporary moral culture. Before the existence of ethical disputes can be cited as evidence against ethical objectivism, one should at least expect some measure for evaluating the scope of these disagreements within and across cultures.

Perhaps, though, the anti-objectivist does not need to offer such a measure, because *any* level of disagreement constitutes a challenge to the objectivity of ethics. But this suggestion smacks of desperate double standards, for no one can seriously suggest that the existence of some cases of error impugns the objectivity of a domain as a whole. At a time when the School Board of Kansas State has recently reintroduced the teaching of creationism in schools, one suspects that few ethical anti-objectivists would see this as impugning the objectivity of biological science, or the justification of evolutionary theory. The mere fact that there are people who come to erroneous judgments is no argument against objectivism about a discourse.

The appeal to the evidence of supposedly widespread ethical disagreement is really rhetoric intended to frighten the ethical objectivist. The fact that concepts are widely *contested* does not mean that they are *contestable*, except in the purely descriptive sense of being 'capable of being contested'. What might begin to challenge ethical objectivism would be a demonstration that a significant portion of ethical concepts were essentially contestable, in the sense of 'deserving to be contested'. The supposed fact of widespread disagreement does little, if anything, to establish that ethical concepts are essentially contestable in this sense. The suggestion might be that it is simply implausible for ethical objectivists to claim, in the face of widespread disagreement about the use of an ethical

concept, that one side is right and the other wrong. But there seems to be plenty of scope for objectivists to refine their position to blunt this challenge. Objectivists can argue that much of the disagreement about the application of ethical concepts can be explained by the parties holding different non-moral beliefs that are intimately involved with the ways in which they deploy the disputed ethical concepts. Particularly in the case of cross-cultural disagreements, the objectivist can argue that differences in situation and context have hidden the fact that the same concept is being applied by both parties.[31]

Nor are objectivists forced to hold the naive view that one party must be right and the other wrong. They can point out that both parties may have an inadequate grasp of the concepts involved, and that, indeed, at this point in time there may be no proper understanding of the concept. This need not be seen as pessimism, but as a reasonable and pragmatic acceptance of the limitations and fallibility of the cognitive and affective powers of human beings. The objectivist can also point to a history of moral development in which concepts that were once hotly contested have ceased to be so as our understanding has developed. Other discourses, which the ethical anti-objectivist or relativist would regard as admitting of objectivity, display a similar history of dispute and contestation in the use of certain concepts, which have since ceased to be contested.[32] And it will be precisely these islands of contention to which we find our attention attracted, rather than the vast oceans of tedious consensus that lap round us.

---

[31]For an example of this sort of defence of moral realism, see Kolnai (1969-70).

[32]Of course, this will be grist to the mill of those who want to be irrealists or relativists of a global kind, for they will see this as supporting the non-objectivity of these supposedly favoured discourses. But this kind of irrealism would hardly help someone seeking to defend an anti-objectivism restricted to ethics.

Admittedly, there does not seem to have been as much progress in ethics as there has been in our understanding of other areas, such as the natural, mathematical and social sciences. But objectivists can admit this and point out that ethics deals with that which is most central to our concerns and interests. It is hardly surprising that our reasoning in this area is particularly prone to being clouded by our self-interest and partiality, not to mention our religious and political views. Again, if anti-objectivists want to turn this pessimism into an argument against objectivism, they have to give a reasoned case for the amount of pessimism that they consider to be incompatible with objectivism or realism. Even the clear fact that ethical concepts seem to exhibit much less determinacy of sense than do concepts concerned with physical science or talk about medium sized dry goods, does not constitute a reason why the contestability of ethical concepts should be *essential*.[33]

Ultimately, these kinds of consideration cannot constitute an a priori argument against objectivism. To achieve this, the anti-objectivist has to go to the heart of the issue and argue that there are simply a priori grounds for thinking that there cannot be any such thing as a cognitive shortcoming in reaching a ethical judgment. Citing evidence of widespread, and apparently unresolvable, disagreement cannot, I have argued, address this need. One way to do this would be to identify the conceptual features of essentially contestable concepts and then to show (probably, not very hard to do) that ethics is marked by a high proportion of these. W.B. Gallie has argued that essentially contested concepts are appraisive concepts, marked by a high degree of internal complexity. This complexity has the consequence that different parties can reach different judgments by

---

[33]For a useful discussion of these issues in the context of aesthetics that sheds light on the topic discussed here, see Sibley (1968). For a more sceptical response to Sibley, see Tanner (1968).

stressing the significance of different elements of the concepts' internal structure.[34] But it is not clear how this proposal might help the ethical anti-objectivist. The *appraisive* nature of the concepts cannot be seen to play into the hands of anti-objectivists without begging the question in their favour, and other areas of thought are marked by internally complex concepts without peril to their objectivity.

In fact, anti-objectivists rarely seem to bother to go to the trouble of offering an a priori characterization of essentially contestable concepts. Wright seems simply to rest his case on an appeal to our intuitions in the face of widespread actual disagreement, while Mackie simply takes the matter for granted:

> Disagreement on questions in history or biology or cosmology does not show that there are no objective issues in these fields.... Disagreement about moral codes seems to reflect people's adherence to and participation in different ways of life. The causal connection seems to be mainly that way round: it is that people approve of monogamy because they participate in a monogamous way of life....[35]

What seems to be going on is that anti-objectivists see the essential contestability of ethical concepts as a priori because they take it that there is no rationally justified basis for forming an ethical judgment. But this assumption is exactly what is at issue in the debate about the conditions on objectivity-sustaining convergence in judgment.

Thus, despite its popularity as an argument against ethical objectivism, it seems that the appeal to considerations about moral disagreement must be regarded either as a posteriori and hopelessly under-specified, or else as a priori, but entirely derivative from, and secondary to, the anti-objectivist's claim (to which I turn next) that ethical judgment cannot meet objectivity-sustaining standards of judgmental convergence. This is not the last that we will hear of issues connected with relativism, for we will find that an account

---

[34]See Gallie (1955-6).

[35]Mackie (1977), p.36.

of objectivity is inadequate if it too easily allows groups making different judgments to count as belonging to their own, hermetically sealed, normative practices. But that issue will be approached from within a discussion of the appropriate constraints on rational convergence and not from any concern with the supposed datum of mass disagreement.

7. Judgmental Convergence and Explanation

Wright's claim that an objectivist is committed to a strong requirement of correspondence to, or representation of, an independently existing reality might seem to be an expression of an idea that is important to any objectivism or realism worth the name.[36] It might be hoped that the fact that judging subjects converged on a (strongly) true judgment could be explained by saying that the convergence was the result of their forming beliefs that corresponded to some independent reality, in the manner of Wright's key making an impression in the wax. Such a convergence in judgment provides a particularly clear way in which judgment could be responsive to something beyond the practice of those judging. Now, earlier in this discussion, we noted that adherence to some correspondence conception of truth is not typical of moral cognitivists or realists, nor indeed realists in general. So we need to look at the motivation that lies behind the claim that correspondence is required for realism and objectivism. We then need to try to express this concern in a way that does not involve an appeal to a contentious notion of correspondence. The debate ought to be set up in such a way that the standard is not so strict that ethical objectivists (or realists in general) can opt out of meeting the demand by

---

[36]An objectivist who does not espouse a full-blooded realism about a domain will still want to maintain that correct judgment matches up to something beyond the thinker, or even the practice to which the thinker belongs. This might, however, be a rational norm (for example, in mathematics), rather than a property of something.

disputing that it was ever something to which they were committed. Nor must the standard be so weak that ethical objectivists can win their spurs too easily. We need to examine the question of what constraints on convergence in judgment must be met in order for a discourse to qualify as exhibiting objectivity in its judgments.

The possibility of convergence in judgment may be necessary for the objective correctness of a judgment, but it is not sufficient. Simple convergence in belief that p among a group of subjects is not enough to guarantee that a belief in p is objectively warranted, since the convergence could be the result of mere collusion. But what is the extra ingredient that is required? One thought, which Gilbert Harman offers us, is that the judgments must not be wholly explicable in terms of the psychological states of the judging subjects, without any reference being made to the properties that the subjects' judgments concern.[37] Thus, Harman claims that cases of 'moral observation' do not meet this standard, because a person's reaching a judgment that, say, torturing cats is wrong, can be explained purely in terms of that person's upbringing and psychology. In contrast, the explanation of a scientist's judging that there is a proton in a cloud chamber makes essential reference to the existence of the proton.

But this way of setting up the requirement is too hasty, because it is far from clear that ethical judgment actually fails to meet it. David Wiggins suggests two reasons for thinking this.[38] First, he points out that if we focus, not on the general judgment that torturing cats is 'wrong', but on the particular reason why it is wrong, namely, 'because it is callously cruel', we want to ask why we were trained to respond in this way. For Harman's case to be sustained, the answer must *not* come back that we are trained to

---

[37]Harman (1977), Ch.1.

[38]Wiggins, "Truth, and Truth as Predicated of Moral Judgments", in Wiggins (1998), pp.156-60.

respond to something that is objectively *there* to be found by those suitably placed and suitably conceptually and perceptually equipped.[39]

The second point that Wiggins makes is that, while Harman says that the moral judgment can be fully explained without reference to anything other than the subject's psychology, the question arises of how this psychological theory is to be constructed.[40] It will have to be a theory that classifies and connects situations, actions, emotional responses and decisions in exactly the way that the subject's (or worse still, the subjects') moral sensibility does; it will have to yield counterfactuals about which judgments or actions will be produced in which situations; and it will have to do all of this without making any essential reference to moral properties. If the psychologist achieves this goal, then, Wiggins suggests, he or she is badly placed to deny substantive truth to the ethical judgments that are now mirrored by his or her theory. Alternatively, the psychologist might try to construct scientific explanations that do not feature properties that connect all instances of, say, cruelty, but that still support counterfactual connections between situations and actions or utterances of evaluative sentences. However, then one is faced with "the sheer unimaginability of someone's supplanting by clinical modes of explanation most if not all of our everyday modes of rendering actions and beliefs intelligible".[41]

A positive picture of the connection between objectivity and judgmental convergence is provided by Wiggins's second "mark of truth":

---

[39]Wiggins (1998), p.157.

[40]Wiggins (1998), pp.157-8.

[41]Wiggins (1998), pp.159-60.

> If x is true, then x will under favourable circumstances command cognitive convergence, and the best explanation of this convergence will either require the actual truth of x or be inconsistent with the denial of x.[42]

Wiggins's suggestion is that convergence in belief that p is of the appropriate sort if the explanation of the convergence in belief has to make appeal to the fact that p. A judgment that p is objectively warranted if one comes to judge that p because there is nothing else to think but that p.[43]

This way of treating convergence in judgment would seem to promise objectivism, without requiring any sort of contentious correspondence. However, one immediately wants to ask, In what sense is there nothing else to think than that p? Presumably, "there is nothing else to think" suggests a normative constraint, rather than one of physical or psychological possibility. It is physically and psychologically possible to think that torturing cats is not cruel, so the thought must be that, within the norms of moral or rational practice, there is nothing else to think but that torturing cats is cruel.[44] The question of whether a particular discourse admits of judgments with some claim to objectivity will then amount to a question of whether the discourse admits of such '*vindicatory explanations*' of why subjects reach these judgments and why they converge in so doing.[45]

But now the worry will be, Why should one respect the norms of this practice rather than those of another? The extent to which Wiggins's proposal succeeds in plotting

---

[42]Wiggins (1998), p.147.

[43]See Wiggins (1998), p.153.

[44]See Moore (1996), for a discussion of how one might refine and extend the application of the idea of 'there being nothing else to...'.

[45]For more on 'vindicatory explanations', see Wiggins (1990-1), pp.66ff.

out a test of objectivity for a set of judgments might seem to be rather limited. Although we have a general schema for what constitutes the right sort of convergence in judgment within a discourse, what counts as *success* in achieving this convergence in judgment might *also* be determined within the discourse.[46] This point can be brought into sharper focus by turning to Wright's suggestion for the further constraint that we need, namely, what he calls, "the width of cosmological role" of the subject matter of the belief. This is:

> [M]easured by the extent to which citing the kinds of states of affairs with which it deals is potentially contributive to the explanation of things *other than*, or *other than via*, our being in attitudinal states which take such states of affairs as object ....[47]

If the facts cited in the explanation of beliefs within a certain discourse also appear in the explanation of beliefs in other discourses, then this indicates that these facts are not just the internal products of the discourse, but pull some genuine explanatory weight. Wright's claim is that ethical properties do not meet this standard of wide cosmological role, because all that they ever explain is the acquisition of moral beliefs and the actions that stem from them.

So, does width of cosmological role provide an insurmountable stumbling block to the ethical objectivist? Is it an appropriate formulation of the further constraint on convergence? An interesting response to this proposal to it comes from Stephen Everson.[48] Suppose that mental events are not identical to physical events and that there is causal interaction between the mental and the physical. Then, Everson argues, a commitment to the width of cosmological role requirement will to force one to the conclusion that there are states within the causal realm that cannot be given an

---

[46]Although, in the end, this worry will be found to be exaggerated (see Chapters 4.4 and 4.5).

[47]Wright (1992), p.196.

[48]Everson (1995), pp.196-9.

explanation. Why so? Suppose we accept Aristotle's view of the acquisition of moral concepts, according to which we acquire the conceptual ability to use these concepts by presentation with instances of those concepts. Then, claims Everson, the only way in which the events that are the actions caused by the application of those concepts can be explained will (in the absence of a reductive account of moral properties) be by appeal to an explanation that unavoidably refers to the moral properties in question. Thus, ethical properties are essentially invoked in an explanation of parts of the causal order, and thus meet the width of cosmological role requirement.

Now, Everson's argument relies on not uncontroversial premises - the truth of an anomalous monism about the mind and Aristotle's conception of the acquisition of evaluative concepts - and so will be as weak or as strong as these premises. However, I wish to highlight two additional worries about this proposal. The first is that it is not clear what work the appeal to Aristotle's views on our acquisition of moral concepts actually does in Everson's argument. Suppose that we concede that there are reaction-independent properties that we respond to in a pre-conceptual way as we come to have an adult set of moral concepts. Presumably these properties count as *moral* ones only because of their relation to a moral theory that is not wholly determined by these properties. If this is correct, then it is not clear how anything on the input side of the belief-action system can help with meeting the width of cosmological role requirement. This, after all, says that, if a putative property explains only our being in certain attitudinal states, then it cannot count as the subject matter of objective judgments. Everson's thought seems to be that we can avoid the accusation that ethical properties are mere shadows cast by our concepts by saying that the acquisition of these concepts involves pre-conceptual encounters with instances of these properties. The problem is how our ethical concepts can hang on to this

good causal pedigree, once we have a fully developed, highly conceptualized ethical understanding that is not wholly determined by these pre-conceptual encounters.

Indeed, it seems open for one to point out that one's understanding of the tasty (a paradigmatically non-objective subject matter) and related concepts could be marked by an acquisition process that is based on pre-conceptual acquaintance with instances of tastiness. These instances will be fairly basic, but, as one develops an adult sense of the tasty, one will acquire more sophisticated conceptions of what is tasty, which a child lacks. The story about a characteristic method of acquiring concepts of the magical or astrological might be harder to explain in this way, but probably not impossible. Children would presumably be presented with certain cases and be told that they are examples of magic at work. They would then go on to develop a way of 'recognizing' future instances of magic. The real question that we need to get at is, Are there properties there for a child to be in contact with in the first place? Therefore, Everson's appeal to the process of acquisition of moral concepts seems to be helpful only if we have already settled this question. But this is precisely the issue that is subject to dispute.

The work in Everson's argument seems to take place on the output side of the belief-action system, with his appeal to considerations about mental causation. This brings me to my second worry, which is that Everson's treatment of the output side would also seem to license objectivism about areas of discourse that we would be very reluctant to regard as objective. Presumably, there are no reductive accounts available for tasty, magical or astrological properties.[49] The concern is that one could mount an analogue of

---

[49]Magical and astrological properties are presumably irreducible because there are simply no such things! The humorous might be another example of such a non-objective domain, but I am increasingly coming to think that discourse about humour shows some surprisingly objective features.

Everson's defence of the reality of ethical properties for any of these three types of properties, because beliefs about what is magical, tasty or written in the stars cause people to do things (for example, respectively: cast spells or attend coven meetings, buy a Big Mac or visit a Michelin Three Star restaurant, and stay in bed because they have had a bad horoscope). So, Everson's attempt to deal with the width of cosmological role requirement seems to render the constraint too weak because of its liberality with what counts as a putative feature of reality having a wide enough cosmological role.

Although Everson's proposal fails, I do think that it inadvertently highlights a better way of dealing with the search for a further constraint on convergence. The basic objectivist intuition seems to be (roughly) that there is something wrong with a conceptual practice that just marches along in step with itself. Appealing to the causal role of something outside the practice might seem to be one way of dealing with this intuition. Certainly, Wright's appeal to the wider causal-explanatory role of the properties to which subjects supposedly respond in forming their beliefs would be a natural way of filling out this idea. But this approach seems to face a dilemma. On the one hand, we can stick to a rigidly scientistic conception of what that role might be, so that legitimate properties have a causal role that is entirely independent of the attitudes of subjects (in a way that Everson's proposal denies). On the other, we can weaken this constraint via an appeal to anomalous monism, but end up with a constraint that is too weak (or so I have argued). Unless Wright gives some account of how discourses are to be individuated, then the width of cosmological role requirement either will be largely empty (since it will not be

clear what constitutes a property playing a role in another discourse), or will have to be so rigidly interpreted that only the properties of physical science pass the test.[50, 51]

We should see Wright's suggestion as just one way of picking up on the intuition that there is something wrong with a practice that simply carries on in its own sweet way. I have sketched out an argument that Wright's requirement must either be empty or else be interpreted very strictly. The strict interpretation, which seems to allow objectivity only to the properties of physical science, is entirely of a piece with the conception of objectivity that has traditionally underpinned ethical anti-objectivism. In the next chapter, I shall argue against the coherence of this conception. If this conception of objectivity is incoherent, then we are free to look elsewhere than a strictly causal-explanatory grounding of objectivity-sustaining convergence. In Chapter 4, particularly in Section 5, I shall continue the discussion of this issue. I shall suggest that, on anything less than a very strictly physicalist interpretation of the width of cosmological role test, there is a case for saying that ethical properties might well pass it. Indeed, it is not clear why objectivists or realists need restrict themselves to a purely causal-explanatory story about width of role.

Ultimately, however, I shall argue that width of role of any kind is, at best, a reassurance of the objectivity of a domain and not a sufficient condition of objectivity or

---

[50]Williamson (1994) develops the worry that Wright's notion of a discourse is so poorly defined that it is not clear that it can do any work. One of the reasons why Everson's reply does not quite work is that, because the notion of a discourse is so ill-defined, it is all too easy to conceive of beliefs or judgments in one discourse feeding through to have an effect in another discourse. Wright responds to Williamson in Wright (1994).

[51]Even if one accepted the strict interpretation of the test, it might still be possible to defend a form of ethical objectivism. Perhaps pain and pleasure pass the test, in which case there are familiar ways in which one might try to understand ethics in terms of these basic building blocks. Since I argue in the next chapter that the conception of objectivity that underlies the test is deeply flawed and I am a little suspicious of the foundational role of pain and pleasure in understanding ethical concepts, I shall not pursue this proposal.

reality. Instead, I shall suggest that the core objectivist intuition can best be satisfied by demanding that judgments within a particular conceptual domain are responsive to reasons that are absolute and are not simply norms specific to the domain.[52] This allows us to develop a test of objectivity that is neutral between physical science, mathematics, ethics, and other putatively objective domains. I shall postpone further discussion of these matters until Chapter 4.

8. Conclusions

In this chapter we have gone some way towards bringing the issues at stake between ethical objectivists and their opponents into focus. To give the opponents of ethical objectivism a good fight I have allowed them to set out the terms of the debate and we have looked at three ways - the error theoretic, the emotivist, and the minimalist anti-realist - in which the anti-objectivist might set up the debate. I have tried to tease out what anti-objectivists find problematic in their opponent's position. The answer seems to boil down to three broad types of problem: issues surrounding queerness and the clash between the prescriptivity of ethical judgments and the supposed motivational inertness of beliefs; appropriate standards of objectivity-sustaining convergence and lack of explanatory power for ethical properties; and, finally, the incompatibility of the supposed essential contestability of ethical concepts and their objectivity. I have argued that, when properly formulated, the last of these problems only provides a priori grounds for challenging ethical objectivism when it rests on a failure of ethical objectivism adequately

---

[52]On one interpretation of Kant, he envisages the Categorical Imperative (CI) as a general constraint on reasons as such, and hence its application to ethical judgments is validated by this wider rational role. Such an interpretation of Kant can be found in O'Neill (1989), Part I.

to account for convergence in judgment. The first group of problems concerning prescriptivity seems to be essentially tied to the non-cognitivist's or anti-objectivist's favoured conception of reality.

This way of setting up the debate makes brutally clear what the objectivist's tasks now are. The first is to undermine the conception of objectivity with which non-cognitivists operate. This is the idea that objectivity is to be grounded only by appeals to a reality conceived of in the perspective-free terms of natural science and that, relatedly, beliefs have to be motivationally inert since their sole function is to reflect that reality. If we can undermine this, then we will have dealt ethical non-cognitivism a hefty, if not mortal, blow. In the next chapter, I shall try to do just this by suggesting that in Wittgenstein's discussion of rule-following and the objectivity of meaning we find a powerful argument against this conception of objectivity. Of course, this still leaves objectivists with a lot to do. Having dethroned the non-cognitivists' conception of objectivity, they still have to show that they themselves have proposed a viable form of objectivity. This is something that I try to do towards the end of Chapter 3. The next task for ethical objectivists is to fill out the conditions on appropriate convergence in judgment that are both genuinely objectivity-sustaining and are more hospitable to ethics than are the causal-explanatory constraints favoured by ethical anti-objectivists. This will be one of the tasks of Chapter 4. Ethical objectivists then need to go on to develop a coherent account of ethical deliberation, in which the judgments reached might meet these standards. They will also need to explain the connection between ethical judgment and action in a manner that does not open up the way for a charge of queerness.

Chapter 3

RULE-FOLLOWING AND OBJECTIVITY

1. Objectivity, Meaning and Rules

So far, I have been slightly wary of using the term 'moral realism' to refer to the meta-ethical view that I shall ultimately try to advance; instead I have preferred to talk about ethical 'objectivism' or 'cognitivism'. However, in this chapter I shall be advocating a general metaphysical picture that, I believe, warrants the label 'realism'. As this might appear to involve a confusion in terminology, I had better explain briefly what is going on. By the end of the last chapter, we had reached the provisional conclusion that underlying many of the ethical non-cognitivists' or anti-objectivists' challenges to their opponents is a conception of what it is for judgments to be objectively correct that is deeply inhospitable to ethical objectivism or cognitivism of any kind. On this view, standards of objective correctness and of reality are utterly independent of us and our interests. My aim in this chapter is to show that Ludwig Wittgenstein provided some compelling arguments against the coherence of this conception, which we might dub 'metaphysical realism'.

Now, a rejection of this position could clearly have welcome consequences for ethical cognitivism or objectivism, because it would open up the possibility of weaker standards of objectivity that ethical judgments might meet. But this weakening of the standards of objectivity might threaten to come at the cost of introducing a general idealism or anti-realism. This might legitimately be regarded as too high a price to pay. Thus, it will be an important aim of this chapter to show that the position we arrive at still

constitutes a form of realism, but one shorn of the fatally incoherent over-ambition of metaphysical realism. So, this chapter is, in part, a defence of realism.

But where does this leave my reluctance to talk of moral realism? I suggested two sorts of reason why I would avoid using the label 'moral realism': it risks confusing associations with scientific, platonic and metaphysical realism, and it is a label associated with a form of ethical cognitivism that sees ethical judgment as a process of moral perception. By the end of this chapter it should be starting to become clear that the first of these worries has been allayed. The sort of realism that I find in Wittgenstein is also hostile to platonic or metaphysical realism. Realism about a domain will be the claim that the judgments within it are capable of being true or false in some substantive, non-minimal sense, but this will involve no claim that there is any metaphysical grounding for truth in correspondence to anything. Thus, on this view, one can be a realist about mathematics if one regards mathematical statements as capable of being true or false, but the truth of mathematical statements is not regarded as a matter of their corresponding to any platonic entities. In this light, there will be little reason to think of realism and cognitivism about ethics as distinct positions. However, there still remains the association of the label 'moral realism' with a perceptual model of ethical judgment, and, as I shall argue in Chapters 5 and 6, this is a position from which I want to distance myself. For this reason, I shall still avoid using the label 'moral realist' for the positive position that I advance.

A second point of clarification concerns the relationship between objectivity and realism. For much of this chapter I shall be discussing realism about meaning. The objectivity of meaning requires that there be objectively correct, 'ratification-independent'

standards for the correct application of concepts.[1] If there are such standards, then we can talk of there being a fact of the matter about what people mean by their use of a word or concept, and thus a realism about meaning is justified. For this reason, I shall treat realism about meaning as equivalent to the claim that meaning displays objectivity. Outside the context of the discussion of the objectivity of, or realism about, meaning, objectivism will not necessarily imply realism. We can conceive of subject matters where there are 'ratification-independent' norms for applying concepts without talk of the truth of judgments (and hence realism) being entirely appropriate.[2] So, it is only in the discussion of the objectivity of meaning that I shall use 'realism' interchangeably with 'objectivism'. These remarks will, at this stage, seem a little sketchy, but I hope that they are clear enough to allay any confusion. With these terminological points of clarification complete, let us now introduce the challenge that this chapter is to discuss.

To put matters very crudely, the truth-value of a sentence depends on two things - the meaning of the sentence and the way the world is. Arguably, one cannot even start to get to grips with the issue of whether or not the sentences of a certain domain of discourse are liable to be true or false, unless there is some objective fact of the matter about what subjects mean by these sentences. In other words, it would seem that an anti-objectivism, or constructivism, about meaning would entail an anti-objectivism or irrealism about facts and truth in general. For this reason, it is possible to approach the issues of realism and objectivism by looking at the objectivity of meaning. Thus, from the point of view of my

---

[1]The phrase 'ratification-independent' is Crispin Wright's (see his essay "Strict Finitism", in Wright (1993), pp.144-7), and means that the patterns of application exist independently of their ever being ratified by the judgments that we might make about them. Also, see McDowell (1984), p.325.

[2]Kant's views on ethics would be the classic example of this. Kant rejects moral realism, but the norms of reason that he appeals to are meant to be thoroughly objective.

wider project, if there is no objectively correct way in which to apply a concept, then objectivism about anything (ethics included) will be in serious trouble. Without there being an objectively correct way to use concepts, it would be hard to regard objectivism of any kind as a live option. Thus, both from the point of view of securing ethical objectivism and of showing that this is compatible with a wider metaphysical picture that is realist rather than idealist, it is essential to see how there could be objectively correct ways of applying concepts.

It seems to be part of our intuitive conception of meaning that, for one to possess a concept, one must be able to deploy that concept in a way that meets genuine, substantive normative constraints on which items count as falling under the concept. Not only that, but one's deployment of the concept must be seen as intentionally aimed at meeting these constraints. The meaning of a concept is, to use John McDowell's phrase, a 'contractual' notion[3] - we commit ourselves by using a word or concept to apply it to the same type of thing in all our future uses of it.

Therefore, the question of what the source is of these substantive normative constraints on meaning and how we commit ourselves to them in grasping a concept will be of the utmost philosophical importance. A common thought is that, in order for this normativity of meaning, and our sensitivity to it, to be explained satisfactorily, it must be explained in terms that are not irreducibly normative. This very broad strategy could take a wide variety of forms. One could explain the normativity of meaning in terms of physical or mental states of the subject (whether these are conceived of as intrinsically intentional

---

[3]McDowell (1984), p.325.

or not). Or one might explain it by seeing the grasp of meaning as a matter of latching on to platonic ideas or features of the world that 'carve it at its joints'.

In this chapter, I want to examine a set of arguments that Wittgenstein developed, which suggest that these ways of understanding the normativity of meaning are deeply mistaken. Wittgenstein couches his discussion of the issue in terms of an examination of what is involved in following a rule. The phenomenon of following a rule acts as a good proxy for meaning, because it, too, involves the idea of acting intentionally in a way that is constrained by a norm. On Wittgenstein's view, the attempt to explain meaning in terms that are not essentially normative or meaning-involving are doomed either to flounder in the mystifying imagery of platonism or else to offer nothing but an irrealism about meaning.[4] A popular interpretation of Wittgenstein's views is that his rejection of platonism, and other forms of metaphysical realism, about meaning suggests that he embraced an irrealism about meaning.[5] This attribution of meaning irrealism to Wittgenstein will be seen to be deeply misconceived and, in any case, a position of dubious coherence. If reductive forms of metaphysical realism about meaning fail, and if irrealism is incoherent, then this might suggest that one should advance a positive account of some form of non-reductive realism. I shall argue that Wittgenstein is best seen as not doing this. Rather, the position that he adopts might be thought of as a form of 'common sense

---

[4]Although Wittgenstein does not use the label 'platonism', it is clear that the conception of rules as ethereal rails has its natural expression in a platonism about mathematics. I do not mean to suggest by the use of this label any view as to whether or not Plato himself really subscribed to such a position. Irrealism about a given discourse is the philosophical view that the predicates that appear in that discourse answer to no real properties. Hence, meaning irrealism is the view that attributions of meaning do not make statements about the real properties of contents.

[5]Kripke (1982) is the most prominent source of this position, but something very similar can be found in Wright (1980).

realism' that rejects the attempt to validate (or, indeed, invalidate) normative practices from some external perspective.[6]

I shall argue that, once the forlorn hope for platonic realism is abandoned, the position that Wittgenstein seems to advocate provides a coherent form of realism. Crucially for my wider project, the view of objectivity and reality that we find in this position is very different from the one that I identified as lying behind the ethical anti-objectivist's opposition to the objectivity of ethical discourse. I shall argue that this conception of objectivity does not seem nearly as obviously incompatible with ethical objectivism.

2. How Not to Follow a Rule (1): Mental States

If one asks how it is possible to come to understand a particular rule, there seems to be an obvious answer - when in doubt one can ask for an equivalent rule that one already understands. This seems to give a perfectly adequate explanation of one's understanding of the rule in a *particular* case. However, the temptation to generalize this strategy, when faced with the question of how it is possible to follow rules at all, is doomed to disappointment, since, clearly, the possibility of substituting one rule for another is exactly the point that is at issue. Now, as I mentioned briefly above, the natural response of many philosophers to this problem will be to claim that, in order adequately to meet this challenge, the answer must involve an appeal to something non-normative. Otherwise, the

---

[6]For more on 'common sense realism' and its attribution to Wittgenstein, see Putnam (1996), pp. 243ff. There seem to be many similarities between this view and the position that John McDowell dubs 'naturalised platonism' (see McDowell (1994)).

normativity of a rule (or, more generally, of meaning) will be left as something *sui generis* and fundamentally unexplained.

Wittgenstein's discussion of following a rule highlights two key features of how we use concepts in a norm-governed or meaningful way. The first is the phenomenological point that we can grasp the meaning of a word 'in a flash'. I can grasp a concept immediately and be utterly certain that I know how to deploy it in any subsequent application. The second intuition is that my grasping a rule involves my committing myself to a pattern of future uses of the concept that is objective and 'ratification-independent'. Any positive account of the normativity of meaning will either have to respect these two intuitions or else need to give a compelling reason why one or other has to be abandoned or revised.

One natural way of approaching the problem is to focus on the state of understanding that a person exhibits when they competently follow a rule. And a natural way of developing this approach would be to try to find the solution to the problem in the mental or physical states of the subject who is following the rule. In this section, I want to look at the first of these two kinds of state - mental states. One might suggest that the normativity of a rule can be explained by a subject's possession of a conscious mental state of understanding, which guarantees correct application of the rule in future cases. If such a state could be identified, the problem of reconciling our two core intuitions might seem to have been resolved, because nothing more complex would be at issue than the possession of a mental state. Since it concerns future applications of a concept, this state would be an example of an intentional state that concerns objects that are not present to the subject. But how is this sort of state to be characterized?

In very broad terms, there are two ways in which one might seek to characterize such an intentional conscious state.[7] On the first view, one might try to characterize the state extrinsically, without any reference being made to the objects that fall under the extension of the concept. On the second view, the characterization would intrinsically involve these objects. The key problem facing the extrinsic conception is this. The objects falling under the extension of the concept are excluded from the initial characterization of the mental state that is supposed to constitute the possession of the concept in question. But it is then hard to see how these objects can ever be satisfactorily seen to be covered by the concept. Wittgenstein pursues this problem from two directions. The first, and less deep, line of attack questions whether there really can be any conscious mental state, possession of which could be regarded as a necessary condition for grasp of a concept. In the sections of *Philosophical Investigations* following *PI* 143, he examines the examples of learning to continue a number series and learning to follow a rule. He draws out the lesson that, not only might there be no one conscious mental state accompanying all cases of understanding a certain concept, but there might, in fact, be no such state at all. Hence, the possession of such a state cannot be a necessary condition for grasping a concept.

Secondly, and much more damagingly, he also argues that there could be no state, possession of which is sufficient to explain grasp of a concept. Wittgenstein's key thought is that, because the state is supposed to be characterized extrinsically, it will be like a sign that is open to multiple interpretations. He introduces the point thus:

> I said that the application of a word is not everywhere bounded by rules. But what does a game look like that is everywhere bounded by rules? Whose rules never let a doubt creep

---

[7]Budd (1984) and (1989), pp.22ff, introduces the distinction in this way.

> in, but stop up all the cracks where it might? - Can't we imagine a rule determining the application of a rule, and a doubt which *it* removes - and so on? (*PI* 84)
>
> A rule stands there like a sign-post. - Does the sign-post leave no doubt open about the way I have to go? Does it shew which direction I am to take when I have passed it; whether along the road or the footpath or cross-country? But where is it said which way I am to follow it; whether in the direction of its finger or (eg.) in the opposite one? (*PI* 85)

The extrinsically characterized conscious mental state cannot determine how one should follow the rule or the sign-post because it too remains open to interpretation. If the interpretation is also extrinsically characterized, then it, too, will stand in need of interpretation, and so on ad infinitum. Wittgenstein sums up this situation at *PI* 198 and 201:

> "But how can a rule shew me what to do at *this* point? Whatever I do is, on some interpretation, in accord with the rule." That is not what we ought to say, but rather: any interpretation still hangs in the air along with what it interprets, and we cannot give it any support. Interpretations come to an end somewhere. (*PI* 198)
>
> This was our paradox: no course of action could be determined by a rule, because every course of action can be made out to accord with the rule. The answer was: if everything can be made out to accord with the rule then it can also be made out to conflict with it. And so there would be neither accord nor conflict here. (*PI* 201)

Thus the key point is that, because of its extrinsic characterization, such a conscious mental state will always remain separated from what the rule requires by a gap that no further interpretation can help to bridge.

3. How Not to Follow a Rule (2): Dispositional States

So, conscious mental states, when extrinsically characterized, will not help to explain how one might grasp a concept and be carried on by that grasp to apply the concept correctly to other objects falling under the concept. But it might be tempting to think that this is simply a problem for the rather naive, and outdated, idea of a conscious state playing this role, since this might seem to smack of an imagism associated with eighteenth century

empiricism. So, it might be thought that a clearly preferable candidate for the state of the subject, possession of which guarantees correct application of the concept, would be some dispositional state of that subject. Here, the thought would be that a subject has grasped a concept if he or she has a disposition to apply the concept correctly. This need involve no dubious phenomenology, and it appeals to ideas that are fully respectable in the naturalistic milieu of late twentieth century analytic philosophy.

If the proposal is something like the claim that I mean by the word 'green' the property to which I am disposed to apply it, then, as it stands, this proposal does not get us very far because, of course, we can be disposed to do many things, including make mistakes. But could this line of thought be developed more successfully so as to account for the normativity of meaning? We might try refining this proposal into something like the claim that I mean by the word 'green' whatever property I am disposed to apply it to when I am in some ideal state. This ideal state will need to be characterized, and, if the reductionist's aspirations are to be met, this will have to be done in terms that are not norm- or meaning-invoking.

Saul Kripke has developed an influential exposition of Wittgenstein's views on rule-following that involves a carefully developed attack on reductive dispositional accounts of meaning.[8] The key point raised against such theories is essentially the same one that we have seen Wittgenstein deploy against an extrinsically characterized conscious mental state of understanding. Such a state, because of its extrinsic characterization, is not capable of determining on its own what does, and does not, accord with the normative

---

[8]See Kripke (1982). My appeal to Kripke's argument should not be seen as a sign that I endorse his overall strategy - I do not. Kripke seriously misunderstands the thrust of Wittgenstein's arguments and the conclusion they are meant to establish. However, a fortunate by-product of this misunderstanding is that Kripke spends a lot of effort developing an attack on dispositional theories of normativity.

constraints of a concept. Summarising Kripke's argument, John McDowell puts the point thus:

> [T]he relation of a disposition to its exercises is in no sense contractual - a disposition is not something to which its exercises are faithful.[9]

Kripke raises two objections to the suggestion that what a speaker means is constituted by some fact about a reductively-defined dispositional state of a speaker. The first problem is that a speaker's meaning or intending something has an infinite capacity - when I use the word 'green' I express a commitment to apply the word to *all* green things, wherever they might be - but dispositions are finite. Ultimately, this problem does not seem to be devastating. On the dispositional account, I need to be disposed to apply the word 'green' even to green objects that are inaccessible. But then the challenge is how one can ever formulate a disposition to apply a word to objects that are, in principle, inaccessible. It is clear, however, that any dispositional account will need to characterize *some* ideal operating conditions for the disposition, and only some idealizations will be genuinely permissible. Although it might be a considerable challenge to the dispositionalist actually to carry perform this task of specifying such permissible idealizations, Kripke does not seem to provide a compelling argument that this cannot, in principle, be done.

Kripke's second argument is that a dispositional theory cannot successfully account for the normativity of meaning in a way that is genuinely reductive. As with the attempt to explain meaning in terms of the possession of conscious mental states, it is once again the extrinsic characterization of such states that means that they will face the problem that Wittgenstein raises at *PI* 86. A dispositional state will simply stand there like

---

[9]McDowell (1984), p.329.

a sign-post offering no guidance when considered on its own, in isolation from an interpretation. Kripke puts the point thus:

> The dispositionalist gives a descriptive account of [the '+'] relation: if '+' meant addition, then I will answer '125'. But this is not the right account of the relation, which is normative, not descriptive ... Computational error, finiteness of my capacity, and other disturbing factors may lead me not to be disposed to respond as I should, but if so, I have not acted in accordance with my intentions. The relation of meaning and intention to future action is normative, not descriptive.[10]

It is certainly not in dispute that meaning is a normative concept (at least not yet). However, the real question is why Kripke might think that pointing out the normative nature of meaning should challenge the idea that no descriptive account of meaning can genuinely capture its essentially normative nature. Does not the dispositionalist have room to develop a more sophisticated account than the crudely descriptive proposal introduced above?

If dispositionalists crudely identify the meaning of a word with all of a subject's dispositions to apply the word, then the result will be deeply counterintuitive in two respects. First, the subject will be disposed to make mistakes - say, applying the word 'green' to blue objects when the light is poor. This means that the dispositionalist account gets wrong the extension of the concept *green*, for the word 'green' now seems to apply to the property *green or blue*. Relatedly, but even worse than this, the account seems to rule out the possibility of error, because subjects just mean whatever it is to which they are disposed to apply a concept. What the dispositionalist needs to do is to identify a select class of dispositions that are genuinely extension-tracking for the concept 'green'. If these dispositions can be seen to result from the proper functioning of a subject's

---

[10]Kripke (1982), p.37.

cognitive mechanisms, then this would suggest that the dispositionalist can give an account of the 'contractual' nature of meaning. But can this task be accomplished?

Paul Boghossian has suggested two conditions on an adequate dispositional theory of meaning. The first is the condition that "any such theory must specify, without presupposing any semantic or intentional materials, property M"[11], where M is the property, possession of which is necessary and sufficient for a subject to be disposed to apply an expression in accordance with its correctness conditions. This condition ensures that the theory offers *extensional correctness*. Boghossian's second condition on a satisfactory dispositional theory is that "it must show how possession of an M-disposition could amount to something deserving to be called a *correctness* condition, something that we would be inherently motivated to satisfy".[12] Two broad categories of candidate have been put forward for the set of dispositions that are meaning-determining: communitarian accounts, which seek to define these dispositions as those displayed by a speaker's linguistic community, and optimal disposition accounts, which try to specify the requisite dispositions by reference to their optimal operating conditions.

The communitarian account is the one that is most interesting from our point of view, because it has been a popular (and, I believe, mistaken) view that Wittgenstein himself endorsed such an account. Crispin Wright puts the communitarian proposal thus:

> None of us unilaterally can make sense of the idea of correct employment of language save by reference to the authority of securable communal assent on the matter; and for the community itself there is no authority, so no standard to meet.[13]

---

[11]Boghossian (1989), p.533.

[12]Boghossian (1989), p.533.

[13]Wright (1980), p.220. He later came to abandon this view (see Wright (1984) and (1986)). For a similar view see Peacocke (1981) and (1984).

Individual speakers have a standard of correctness because their dispositions to use words can be checked against the dispositions of their entire linguistic community. The community's dispositions have no check on their correctness. The community does not go right or wrong, "rather, it just goes".[14] It is clear that this theory rules out the apparently coherent thought that there might be situations in which an entire linguistic community goes wrong in its use of a term. Intuitively, it seems to make sense to think that, if we all woke up tomorrow and started classifying together red and blue objects and calling them both 'red', we would be making a systematic mistake. But the communitarian account seems to deny that we would be making such a mistake. Instead, this would just be how the word 'red' is to be applied. It is pretty clear that the communitarian account, at best, offers a heavily watered down conception of the objectivity that we take meaning to involve. I shall delay until Section 5 a full discussion of the coherence of this position.

The other variety of reductive dispositionalist account claims that the meaning-determining disposition to apply a word is the disposition to apply it in certain specifically defined conditions in which subjects are deemed to be incapable of error. Thus the extension of an expression can be identified with whatever property it is to which subjects are disposed to apply the expression in those conditions. For our purposes, it would take us too far off course to consider the various proposals for how to specify what these special conditions are. However, it is still possible to give such theories a little more consideration than Kripke does when he dismisses them out of hand.

The challenge facing reductive dispositionalists is to give a purely reductive specification of the optimality conditions in non-normative, non-meaning-involving terms,

---

[14] Wright (1980), p.220.

and the disposition thus specified must lead the subject to apply a term only to what it really means. The problem that the dispositionalist faces is the holistic nature of belief. Donald Davidson puts the point thus:

> There is no assigning beliefs to a person one by one on the basis of his verbal behaviour, his choices, or other local signs no matter how plain and evident, for we make sense of particular beliefs only as they cohere with other beliefs, with preferences, with intentions, hopes, fears, expectations, and the rest.[15]

If our beliefs were assignable one by one, then the dispositionalist might hope to specify optimality conditions one by one for a subject holding each belief type. However, this hope is dashed by the holism of belief (and of our other propositional attitudes). In order to specify the optimal conditions under which I am disposed to think 'There is a cat on the mat', the dispositionalist will have to specify the optimal conditions for a host of other background beliefs. It is not enough to say that the optimal condition is there being a cat on the mat, for I also need to believe that I am not drugged, or dreaming, or subject to hallucinations, or surrounded by small dogs disguised as cats or..... The list of these other beliefs that need to be specified is endless. An optimality condition would have to be such as to exclude all of this potential infinity of beliefs that could make me believe that there is a cat on the mat, even when there is not. The dispositionalist needs to pull off the task of giving this specification in entirely non-normative and non-meaning-involving terms. This task looks to be impossible since, as Davidson puts it, there "is no echo in physical theory" for the norms of rationality that specify the conditions under which we believe something.[16]

---

[15]Davidson (1982a), p.221.

[16]Davidson (1982a), p.231.

Holism, then, means that there is no entry point for the dispositionalist to get into our system of beliefs and start building a reductive account of their optimality conditions. And, of course, if there can be no account in these terms of when subjects are in an optimal condition to belief that p, there can be no reductively defined optimal conditions for their correctly asserting that p.[17]

What we have seen, then, is that Wittgenstein's challenge at *PI* 86 is not merely an attack on some rather quaint and outmoded idea that meaning might be constituted by possession of certain conscious mental states. Rather, it is an attack on the very idea that meaning can be given a characterization in terms that are extrinsic or reductionist. This means that even sophisticated forms of reductionist dispositionalism about meaning also seem vulnerable to the same challenge.

4. How Not to Follow a Rule (3): Platonism

We have so far been examining the extrinsic characterization of states of meaning because of the clear affinity of such accounts with one of the central projects in contemporary philosophy, namely reductionism. But such approaches seem to force us into the situation that Wittgenstein highlights when he writes:

> [I]n the course of our argument we give one interpretation after another; as if each contented us at least for a moment, until we thought of another standing behind it. (*PI* 201)

Because extrinsic characterizations (whether in terms of conscious mental states or dispositions) do not make reference to the objects falling under the extension of the

---

[17]For more discussion of the implications of holism for the irreducibility of the mental, see various essays in Davidson (1982a), particularly "Mental Events" and "Psychology as Philosophy", and Child (1993) and (1994), Ch.2.

concept being characterized, they seem to demand further interpretation. But these further interpretations suffer from the same defect, leading us into an infinite regress.

The obvious response to this challenge is to seek a characterization of meaningful states that is intrinsic in some way. On this view, the mental state of meaning a word is only characterizable by reference to the future applications of the word. This offers the hope that possession of such a state would not leave a gap between meaning and application, and hence there would be no need for any further interpretation. Wittgenstein outlines this idea at *PI* 188:

> Here I should first of all like to say: your idea was that that act of meaning the order had in its own way already traversed all those steps: that when you meant it your mind as it were flew ahead and took all the steps before you physically arrived at this or that one. Thus you are inclined to use such expressions as: "The steps are really already taken, even before I take them in writing or orally or in thought." And it seemed as if they were in some unique way predetermined, anticipated - as only the act of meaning can anticipate reality.

Despite anything that Wittgenstein's interlocutor might want to say, this conception simply takes us out of the frying pan of the endless regress of interpretations and lands us in the fire of mystery and platonic mythology. What sort of state is this mental state of grasping a meaning? It clearly cannot be an ordinary conscious mental state, because no conscious state could represent an infinite array of objects to which the concept is to be applied. Instead, the state must be one that will avoid the demand for constant reinterpretation:

> What one wants to say is: "Every sign is capable of interpretation; but the meaning mustn't be capable of interpretation." (*BB*, p.34)

This 'last interpretation' must be able to bridge the gap between meaning and application, and must resist the sceptical demand for further interpretation. Wittgenstein sees this demand as suggesting a picture of the grasp of meaning as a latching on to an ethereal set of rails that sweep the speaker on to correct application of the concept in all future cases:

> Whence the idea that the beginning of a series is a visible section of rails invisibly laid to infinity? Well, we might imagine the rails instead of a rule. And infinitely long rails correspond to the unlimited application of the rule. (*PI* 218)

This conception of the normative force of a rule is simply mythology, and for two reasons it cannot provide a coherent basis for a theory of meaning. The first reason for this is that, when the rules are conceived as being like rails that automatically take us from an initial instruction to all future applications, they must be quite unlike normal rails. In order to avoid the re-emergence of the problems faced by the extrinsic characterization, they must carry us on to future applications without any mishap, and thus must be regarded as, to use McDowell's phrase, 'super-rigid'.[18] Unlike normal rails, which can bend and distort, these super-rigid rails do not admit of any such corruption, since their very purpose is to give an absolutely secure guarantee of correct application of a concept.

Wittgenstein raises this point in his discussion of 'the machine as symbolizing its action' at *PI* 193 and *Z* 296:

> How queer: It looks as if a physical (mechanical) form of guidance could misfire and let in something unforeseen, but not a rule! As if the rule were, so to speak, the only reliable form of guidance.

Wittgenstein argues that, because an ordinary machine, or mechanical form of guidance, can be distorted to produce movements unintended by its designer, the platonist model of a mechanism must be of a radically different sort. The platonist is forced to regard the correct movements of the machine as present only in the 'machine-as-symbol':

> But when we reflect that the machine could also move differently it may look as if the way it moves must be contained in the machine-as-symbol far more determinately than in the actual machine... And it is quite true: the movement of the machine-as-symbol is predetermined in a different sense from that in which the movement of any given actual machine is predetermined. (*PI* 193)

---

[18]McDowell (1984), p.332.

The platonist wants to model our meaning something by an expression on the movement of a machine. But reflection shows that in any real machine distortion is possible, so platonists must say that the machine they have in mind - the 'machine-as-symbol' - is harder, more rigid than any actual machine in that its movements are absolutely guaranteed. This rigidity is what creates the profound difference between a real mechanism and the 'machine-as-symbol', yet it was this very rigidity that the whole analogy with the workings of a machine was meant to explain! So Wittgenstein concludes:

> You have no model of this superlative fact, but you are seduced into using a super-expression. (It might be called a philosophical superlative.) (*PI* 192)

Ultimately, then, platonists cannot cash out the central metaphor that their view rests upon, without appealing to the very 'superlative fact' that they are trying to convey.

If the first objection to platonism about rules is that it is conceptually incoherent, the second is that there is no plausible picture of what it would be for a subject to grasp such a rule. The epistemology of the relationship of a subject to a rule-as-rail is totally mysterious. Wittgenstein has his platonist interlocutor say:

> "But do you really explain to the other person what you yourself understand? Don't you get him to guess the essential thing? You give him examples, - but he has to guess their drift, to guess your intention." (*PI* 210)

The platonist seems to rely on an epistemology of guessing. As Crispin Wright puts it, this is "the idea that understanding an expression is a kind of "cottoning on"; that is, a leap, an inspired guess at the pattern of application which the instructor is trying to get across".[19] At the start of Section 2 I introduced two core intuitions about meaning, the second of which was that one can grasp the meaning of a concept utterly firmly and in a

---

[19]Wright (1980), p.216.

flash. This idea of grasping a platonic rail as a fortuitous 'cottoning on' simply fails to do any justice to this intuition.

5. The Incoherence of Meaning Irrealism

The preceding discussion seems to have left us at an impasse. If we try to characterize states the possession of which constitute a grasp of a concept in extrinsic terms that make no essential reference to the objects that fall under the concept, then we seem to end up in an infinite regress of interpretations. This seems to lead to the conclusion that there can be no meaning:

> This was our paradox: no course of action could be determined by a rule, because every course of action can be made out to accord with the rule. The answer was: if everything can be made out to accord with the rule, then it can also be made out to conflict with it. And so there would be neither accord nor conflict here. (*PI* 201)

Without the notions of accord and conflict with a norm, there can be no room for correctness, and hence no meaning. On the other hand, if one rejects the extrinsic form of characterization, with its demand for endless interpretations, one is faced with the bankrupt mythology of platonism about rules. Where do we go from here?

In this section, I shall consider, and reject, an influential response to this sorry state of affairs. On this view, Wittgenstein is seen as accepting the first horn of this dilemma, namely the conclusion that there are no objective facts about meaning. Not only is this position of meaning irrealism dubiously coherent, there is every reason to think that Wittgenstein did not endorse it.

It is one of the key features of Saul Kripke's interpretation of Wittgenstein's views on rule-following that it casts him as accepting a 'Sceptical Paradox'. This is the thesis that there can be no fact about a subject that would constitute his or her meaning one

thing, rather than another, by the words that he or she uses.[20] Kripke thinks that Wittgenstein then goes on to appeal to communal norms of assertibility as a 'Sceptical Solution' to help us to live with the otherwise "insane and intolerable" situation of irrealism about meaning. Before examining the plausibility of this projectivist, non-factual approach to meaning, I want to question the coherence of the position that Kripke sees Wittgenstein as accepting, namely irrealism about meaning.

Faced with the discovery that factualism about a discourse is unsustainable, one might respond by suggesting an error theory for that discourse, according to which statements in the discourse attempt to state facts, but, in fact, turn out to be universally false. Paul Boghossian has suggested a simple argument as to why this proposal will not work in the case of irrealism about meaning. An error theory of meaning would state something like:

(1) For any S: 'S means that p' is false.

The disquotational nature of the truth predicate means that (1) entails:

(2) For any S: 'S' has no meaning.

Because (2) entails that no sentence can have a meaning, (1) itself is not true because it says that some sentences, namely those attributing meaning, are false and hence meaningful.[21]

This leaves us with the other approach open to one who adopts irrealism about a discourse, namely non-factualism. As has been noted by several commentators, non-factualism about meaning has the unusual property of automatically globalizing. Non-

---

[20]See Kripke (1982), pp.66ff.

[21]Boghossian (1989), p.523, and, for a more extensive treatment of the argument, Boghossian (1990), section III.

factualism about meaning is the denial that any meaning-attributing statement is truth-conditional. Hence, because the truth-conditions of any sentence will, at least in part, be a function of its meaning, we are driven to a non-factualism about truth-conditions themselves.[22] Crispin Wright has suggested that this tendency of non-factualism about meaning to globalize is the source of a deep incoherence in the position. This is because, unlike other forms of non-factualism, it would not be able to point to a class of fact-stating sentences that are to be contrasted with the ones that it has discovered to be non-factual, since its globalization means that there are *no* fact-stating sentences. Whilst this characteristic of non-factualism about meaning is certainly remarkable, it is not clear that the position is thereby actually incoherent.

Paul Boghossian, however, provides another argument against the coherence of such a view. He argues that non-factualism about meaning requires a contradictory conception of truth. Non-factualism about a discourse must be seen as the denial of a substantial, more than purely minimal, truth to sentences of the target discourse. Boghossian glosses this as the claim that "the predicate 'true' stands for some sort of language-independent property, eligibility for which will not be certified purely by the fact that a sentence is declarative and significant".[23] Now, the question of whether an object possesses a 'robust' property such as this will have to be a factual one. Thus, if truth is robust (as non-factualism's denial of truth to sentences in the target discourse requires), then judgments about the truth-value of a sentence will have to be factual. But, as we have just seen, the globalization of non-factualism about meaning denies precisely this claim.

---

[22]See, for example, Boghossian (1989), p.524, and Wright (1984), pp.769-70.

[23]Boghossian (1989), p.526.

This yields a straight contradiction at the heart of non-factualism about meaning because it requires truth to be both 'robust' and minimal.[24]

These arguments certainly add weight to the suspicion that there is something strange and paradoxical about meaning irrealism, but it would take us too far out of our way to engage in any more detailed assessment of their cogency. Instead, I want to move on to consider the plausibility of the coping strategies deployed by those who accept the non-factuality of meaning. If, as I shall argue, these fail to convince, then one who takes the non-factualist route is forced to face up to all the undiluted consequences of that position.

Kripke sees Wittgenstein as trying to soften the bitter pill of the Sceptical Paradox by offering a Sceptical Solution to the paradox, which appeals to conditions of communal assent to allow for the existence of purely descriptive assertibility conditions.[25] In this way, it is hoped to salvage some notion of correctness of meaning, even if this does not have the objectivity that we intuitively (but incorrectly, if the Sceptical Argument is to be believed) attribute to meaning. The conditions have to be assertibility-conditions, rather than truth-conditions, otherwise the Sceptical Argument would simply re-emerge to challenge their coherence. In order for these conditions of justified or warranted

---

[24]Wright, (in Wright (1992), appendix to Ch.6), disputes Boghossian's conclusion, because, in his view, the 'meaning-minimalist' need only claim that there is a *correct* answer as to whether meaning-minimalism is to be accepted. Thus there is no straight contradiction of the sort that Boghossian claims to have unearthed. However, Wright does admit that if the acceptability of meaning-minimalism is only correctness-apt, then the answer to the question of its acceptability would not exhibit Cognitive Command. Thus, the view lacks "a certain kind of cogency", because one could refuse to accept it without being guilty of a cognitive shortcoming.

[25]I shall not go into the exegetical case against the attribution of such a view to Wittgenstein. My own suggestions about how to read Wittgenstein's positive position are given below in Sections 6 and 7, and it should be fairly clear that this interpretation is radically different from Kripke's. The cogency of Kripke's exegesis is questioned by, amongst others, McGinn (1984), Baker and Hacker (1984), McDowell (1984), and Putnam (1996), pp.252-60.

assertibility to offer anything resembling normativity, they must allow for a distinction between actually following a rule and merely seeming to do so. Here, Kripke brings in the notion of communal assent to supply the requisite distinction. There can be a genuine distinction between my merely thinking I am following a rule and my actually doing so, because I only count as doing the latter if my actions are in line with the actual dispositions of my linguistic community. Kripke's assertibility conditions can be couched only in descriptive terms about how the community is disposed to ascribe meanings and assess speakers' responses.[26]

It should be clear that this position seems very similar to the communitarian attempt to define meaning-giving dispositions, outlined in Section 3 above. That view would say that there was a fact of the matter about whether my linguistic behaviour corresponded with the dispositions of my linguistic community. However, the extent to which this would be a fact in any substantial sense is questionable, since the communitarian denies that there is any sense in which the community itself responds in its use of a word to an independently existing meaning. Because of their obvious similarities I shall tackle the two views together.

There are three problems that I want to raise about the attempt to explain normativity in terms of communal assent. The first problem concerns the descriptive

---

[26]It is a well known feature of Kripke's position that he sees this as automatically precluding the possibility of a private language, since, in the private linguist's case, the assertibility conditions will be empty. Some have argued that it turns out that Kripke's communal assertibility conditions can, in fact, only be formulated in a way that is parasitic on solitary conditions of assertibility (see Goldfarb (1985), McGinn (1984), and Boghossian (1989), pp.520-2). On this view, Kripke gets the issue completely the wrong way round. It is also not clear that a private or solitary linguist cannot deploy a similar manoeuvre by appealing to stratified dispositions (see Blackburn (1984a)).

nature of the conditions of communal assent.[27] The theory is very modest in its aspirations - the assertibility-conditions merely record in descriptive terms those conditions under which the community actually endorses a speaker's responses. This brings into sharp focus the way in which the theory has to pass over any attempt to meet the intuitive demand that meaning has a 'contractual' aspect.[28]

This leads naturally on to our second challenge to the communitarian picture. This is the objection that conditions of communal assent just do not seem to offer anything remotely like the sort of objectivity that we associate with meaning. One common way of bringing this out is to point to the fact that there seems to be no incoherence in the idea that a linguistic community could non-collusively go wrong in its use of a term, and yet the communitarian account rules out this possibility by definition. Communitarians will respond that this objection simply misunderstands their view - they are not trying to give an account of our intuitive idea of the objectivity of meaning, but are trying to come up with an account that is purged of the dubious baggage of platonism. The question, then, is whether this conception is so etiolated that it fails to offer *any* notion of objectivity worthy of the name. Once we see that really there are no norms, but that the best that can be salvaged is some form of endorsement by the other members of our linguistic community, it then seems to be an act of philosophical bad faith to be satisfied with this substitute. We should, rather, accept the Sceptical Paradox in an unadulterated form and

---

[27]This point is most obviously directed at Kripke's version of the theory, but Wright's is also open to this charge because, on his view, there is nothing underpinning the responses of the community - "it just goes".

[28]It also makes it hard to see how such assertibility-conditions could have the necessary modal force to support the claim that there *could* not be a private language (see Boghossian (1989), p.520).

recognise that a key aspect of our intuitive picture of meaning is hopelessly incoherent. John McDowell puts this point thus:

> But the purging yields the picture of what I have been calling 'the basic level'; and at this level Wright's picture has no room for norms, and hence - given the normativeness of meaning - no room for meaning... It is problematic, however, whether the picture of the basic level, once entertained as such, can be prevented from purporting to contain the real truth about linguistic behaviour. In that case its freedom from norms will preclude our attributing any genuine substance to the etiolated normativeness that Wright hopes to preserve.[29]

The third objection to the communitarian account is that it fails Boghossian's 'extensional correctness' requirement, which was introduced in Section 3 above.[30] Communities, by their very nature, are made up of individuals with similar capacities. This feature of communities means that they are liable to experience systematic errors. If I am likely to mistake horses for cows on a dark night (to use Boghossian's example), so are the other members of my linguistic community. If communal dispositions to apply words are to be the sole determinant of meaning, then the meaning of the word 'horse' will be the disjunctive property *horse or cow*. It would appear that the communitarian theory will give the wrong account of the extension of a term precisely where errors occur for the most basic of reasons (and are thus liable to be systematic in nature). This objection bolsters our concern that the sort of normativity that the communitarian account offers is far, far too etiolated, for it suggests that the community will take the wrong course just when each individual is most likely to do so.

Attempts to salvage some highly weakened form of objectivity from the wreckage left by the acceptance that there are no real norms or facts of meaning fail to deliver anything that might offer a satisfactory account of the objectivity of meaning. So the

---

[29]McDowell (1984), p.336.

[30]Boghossian (1989), p.536.

lesson seems to be that non-factualism about meaning would have to be accepted in all its incoherent entirety, if the arguments for it are sound.

6. Wittgenstein's own Views on Rule-Following

The discussion so far seems to have left us in a deeply perplexing position. Extrinsic characterizations of meaning seemed to lead to an endless regress of interpretations and irrealism about meaning, while the platonist alternative was hopelessly mythological. We then found that meaning irrealism seemed to be of rather dubious coherence, and the attempts to soften this conclusion by appeal to communal standards of correctness did little to address the problem. At this point Wittgenstein's own views on rule-following and meaning should appear more attractive than they otherwise might if bluntly asserted straight out of the blue.

The position that Wittgenstein adopts is a subtle one that can easily be mistaken for one of the alternative positions that we have discussed and rejected. His key insight is that the platonist, the reductionist and the irrealist make a common error: they all assume that our practices of following a rule or using a word in a meaningful way must be validated from an external perspective. The platonist tries to explain the compelling nature of the objectivity of meaning by suggesting that our use of concepts must involve our latching on to rules-as-rails that guarantee correct application of a concept. Logic is then thought of as a kind of ultra-physics that plots out the course of these rails. Reductionists try to explain our grasp of a concept in terms either of the possession of a certain conscious mental state or of our being in a certain dispositional state. They are also trying to give an account of the normativity of meaning from outside our norm-governed practices, without any appeal to normative concepts. They are then forced into trying to

take the perfectly mundane notion of an interpretation and stretch it to a point where it loses all content, a point Wittgenstein makes in the last paragraph of *PI* 201:

> Hence there is an inclination to say: every action according to the rule is an interpretation. But we ought to restrict the term "interpretation" to substitution of one expression of the rule for another.

Interpretation has a perfectly legitimate role as a means of getting from one way of understanding a concept, or rule, to another way of understanding it. However, it cannot provide a miraculous means of transforming brute behaviour, conceived of in non-normative terms, into norm-governed activity. Meaning irrealists share with both platonists and reductionists the same basic conception of what an adequate theory of the normativity of meaning must achieve. They take the failure of both these approaches as compelling us towards the conclusion that the ascription of meaning is not a factual or objective matter.

Wittgenstein rejects these attempts to validate our meaningful use of concepts from a standing start. His rejection of the mythology of platonism and the hopeless regress of interpretations does not amount to an endorsement of meaning irrealism or constructivism about meaning, as a reading of the second paragraph of *PI* 201 makes clear:

> It can be seen that there is a misunderstanding here from the mere fact that in the course of our argument we give one interpretation after another; as if each one contented us at least for a moment, until we thought of yet another standing behind it. What this shews is that there is a way of grasping a rule which is *not* an *interpretation*, but which is exhibited in what we call "obeying the rule" and "going against it" in actual cases.

The 'misunderstanding' is the attempt to explain meaning in non-normative terms, which necessarily fails. But what form does Wittgenstein's alternative account take? What is this 'way of grasping a rule which is not an interpretation'?

Strictly speaking, if we understand the structure of the debate in the way that I have suggested, then Wittgenstein does not need to give much of an answer. The answer as to how one grasps a rule without an interpretation will be provided by giving an account of what would ordinarily count as grasping a rule. Thus, if one asks what makes it correct for one to use the word 'green' to describe the colour of fresh grass, the answer is simply the fact that the grass is green. This will look like a very inadequate answer to the platonist, reductionist or irrealist, but I have argued that the point of Wittgenstein's discussion of rule-following is to show how mistaken it is to seek something that explains "How We Are Able to Follow Rules".[31] Wittgenstein shows that the demand for such a thing is confused and ends up leading to incoherence, however one tries to meet it. His rejection of a platonic realism about rules does not mean that he embraces an anti-realism or constructivism instead. Rather, his is a rejection of a specifically philosophical, or metaphysical, realism that is meant to make a 'common sense realism' about rules seem unproblematic. By the lights of this common sense realism, it is entirely appropriate to offer the answer, 'Because it *is* green', to the question, 'Why is it correct to describe the colour of fresh grass as 'green'?'

I shall say more in the next section about how this approach does succeed in providing a form of objectivity, but in this section I want to say a bit more about the positive features of Wittgenstein's account of rule-following. Again, it is important to stress that he is not attempting to give a constructive answer to something that he recognizes as a genuine question. Rather, observations about features of the practices of following a rule, continuing a number series or using words in a meaningful way are

---

[31]See Putnam (1996), p.245, for this phrase.

offered as 'reminders' (see *PI* 127) that will lessen any attraction one might feel towards a constructive philosophical account of these practices.

In assembling these reminders to show us how we engage in norm-governed activity, Wittgenstein's key idea is that of a practice or custom. His remark, "And hence also 'obeying a rule' is a practice" (*PI* 202), is offered, not as an answer to the question of how we might rid ourselves of the need for interpretation, but as part of a rejection of the notion that this demand is legitimate. A practice is not understood as activity, describable in non-normative terms, that can generate meaningful behaviour out of nothing, in the manner in which Wright and Kripke seem to think that communal endorsement can generate norms of meaning out of brute behaviour. Rather, such a practice is conceived of as being an ordinary, and irreducibly normative, feature of our rational lives. The fact that it is irreducibly normative is unproblematic, since no constructive account of following a rule is being offered or demanded. So what does Wittgenstein mean by a practice?

What seems central to his concerns is the idea of a regular activity into which we are initiated by training, so that our engagement in the practice becomes natural. The idea is expressed at *PI* 198:

> "Then can whatever I do be brought into accord with the rule?" - Let me ask this: what has the expression of a rule - say a sign-post - got to do with my actions? What sort of connection is there here? - Well, perhaps this one: I have been trained to react to this sign in a particular way, and now I do react to it.
>
> "But that is only to give a causal connection; to tell how it has come about that we now go by the sign-post; not what this going-by-the-sign really consists in." On the contrary; I have further indicated that a person goes by a sign-post only so far as there exists a regular use of sign-posts, a custom.

What explains the connection between the rule and one's actions in accordance with it is not the possession of an interpretation but, one's simple reactions that have been developed by the training that one has gone through in becoming a participant in a

practice. When Wittgenstein's interlocutor raises the concern that this is merely to cite a causal connection and does not explain 'going-by-the-sign', Wittgenstein replies that this objection ignores the fact that the training is training into a normative practice or custom. There would be no custom without the unreflective activity to which the training gives rise, but, on the other hand, the training would not count as training in following sign-posts without the existence of the custom. The way that Wittgenstein stresses the interdependence of the training and the normative aspect of the custom suggests that any form of reductive account of the custom in terms of dispositions would be untenable. It also illustrates the way in which he is offering reminders, rather than a constructive account.

The unreflective action that underpins the practice is revealed in each case of following a rule, as Wittgenstein makes clear at *PI* 218:

> When I obey a rule, I do not choose.
> I obey the rule *blindly*.

And again at *OC* 204:

> Giving grounds, however, justifying the evidence, comes to an end;- but the end is not certain propositions' striking us immediately as true, i.e. it is a kind of *seeing* on our part; it is our *acting*, which lies at the bottom of the language-game.

Our norm-guided action is based on simple, non-reflective responses and is not mediated by interpretations or justifications:[32]

> How can he know how he is to continue a pattern by himself - whatever instructions you give him? - Well, how do I know? - If that means "Have I reasons?" the answer is: reasons will soon give out. And then I shall act, without reasons. (*PI* 211)

---

[32]See Winch (1981) for an insightful discussion of the role of action in Wittgenstein's account of how we avoid the need for interpretation.

By saying that one acts without reasons Wittgenstein is not suggesting that norm-guided activity is constructed out of a norm-free 'basic level', to use McDowell's phrase.[33] Instead, his point is that the process of justification is one that takes place within our norm-governed practices, and within these practices the process of asking for justifications cannot go on forever without losing its intelligibility:

> If I have exhausted the justifications I have reached bedrock, and my spade is turned. Then I am inclined to say: "This is simply what I do". (*PI* 217)

The lack of justification "does not trouble me" (*PI* 212) because it is this very 'bedrock' of simple, unreflective action that provides the context within which justifications can be offered. Our inability to justify this behaviour does not leave us open to criticism, because, "To use an expression without justification does not mean to use it without right" (*PI* 289). In particular, there is no need for a metaphysical realist validation of this right. Our recognition of the way in which we act without justification is a recognition of the role of this type of action as part of the fabric of ordinary, norm-governed practices. That this behaviour is both simple and already normative in character is revealed by *RFM* VI-28:

> Following according to the rule is FUNDAMENTAL to our language-game. It characterizes what we call description.

Although the action is fundamental to the language-game, it is not to be described in terms of brute physical goings-on, but in norm-laden terminology as 'following according to the rule'.

Simple, but irreducibly normative action is, then, the first of the features of our rational, norm-governed practices about which Wittgenstein wants to remind us; the second feature that I want to examine is regular activity. Wittgenstein expresses this idea at *PI* 199:

---

[33]McDowell (1984), p.336.

> Is what we call "obeying a rule" something that it would be possible for only *one* man to do, and to do only *once* in his life? - This is of course a note on the grammar of the expression "to obey a rule".
>
> It is not possible that there should have been only one occasion on which someone obeyed 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 obey a rule, to make a report, to give an order, to play a game of chess, are *customs* (uses, institutions).

The context of regular, repeated normatively governed activity is a necessary feature for rational practices to count as genuinely norm-governed. But notice how weak the requirement is. All that Wittgenstein demands is that the activity not be carried out just once, by one person. This is a far-cry from the picture of communal assent as constitutive of normativity that we found in Wright's and Kripke's readings of Wittgenstein.[34] Wittgenstein's refusal to endorse a communitarian view is repeated again at *RFM* III-67:

> 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 once in his life.

One would have thought that, if Wittgenstein is espousing a communitarian view, he would have no reason to hesitate in answering his question with a resounding 'Yes'.

Wittgenstein's reluctance to say anything particularly detailed about how much regularity is required for a normative practice to exist is entirely explicable if, as I have argued, he is not suggesting a constructive theory, but merely offering reminders about the features of our actual practices. He does not say that every rule must be followed at least once (let alone be followed by a community), rather his point is that rule-governed activity must take place within a wider context of regular activities.

---

[34]There is a large debate about the precise relationship between community and custom in Wittgenstein's discussion, which I shall not even try to survey. Baker and Hacker (1984) and (1989), McGinn (1984), and Budd (1984) and (1989), have argued that Wittgenstein's references to practices and customs make no essential appeal to community. Wright and Kripke, of course, do think that there is such an essential appeal to community (although Wright has subsequently changed his views on this - see footnote 13). McDowell (1981a) and (1984), and Pears (1988), adopt an intermediate position where there is an appeal to community, but not as part of a constitutive account of the normativity of rules or meaning.

A consequence of this reminder is a rejection of an aspiration that both the platonist and the reductionist cherish, namely that one could conceive of what grounds the normativity of a rule - a platonic rule-as-rail or a behavioural disposition - in isolation from any context. Wittgenstein sees the observation of our ordinary practices as revealing that a number of different features of a situation - including training and regularity of activity - can determine whether an activity is rule-governed. At *RFM* VI-35 the importance of regularity is stressed when Wittgenstein discusses what we ought to say about a reproduction of England that lasts for only two minutes. The people in this country might appear to be engaged in the same sort of rational, norm-governed activity as people are in the real England. However, Wittgenstein is keen to stress that this assessment might be too swift:

> Ought we to say that the two-minute man is calculating? Could we for example not imagine a past and a continuation of these two minutes, which would make us call the process something quite different? (*RFM* VI-35)

Wittgenstein is not saying that we should definitely not regard the two-minute man as calculating, but that the highly truncated nature of the context in which we view him means that it is not clear that he is calculating. In bringing to our attention the role of regularity in our rule-following activities, Wittgenstein is merely pointing out that talk of rules would make no sense without the existence of the regularities that this talk presupposes.

The final feature of our rule-following activity that Wittgenstein highlights is shared training and the shared responses that the training produces. In order to communicate with others, one must be able to understand them without the need for endless interpretation. We have seen that the possibility of understanding without interpretation depends upon a capacity to act unreflectively and without justification. It

follows directly from this, that, if we are to understand others without interpretation, we must share a common tendency to act unreflectively in the same way. This idea is captured by Wittgenstein at *PI* 241-2:

> "So you are saying that human agreement decides what is true and false?" - It is what human beings *say* that is true and false; and they agree in the *language* they use. That is not agreement in opinions but in a form of life.
>
> If language is to be a means of communication there must be agreement not only in definitions but also (queer as this may sound) in judgements. This seems to abolish logic, but it does not do so.- It is one thing to describe methods of measurement, and another to obtain and state results of measurement. But what we call "measuring" is partly determined by a certain constancy in results of measurement.

In order for there to be language in which contents accessible to more than one subject can be meaningfully expressed, there must be a form of life shared between the parties. This picture is vividly captured by Stanley Cavell when he writes:

> That on the whole we [do project words into novel contexts] is a matter of our sharing routes of interest and feeling, modes of response, senses of humour and of significance and of fulfilment, of what is outrageous, of what is similar to what else, what a rebuke, what forgiveness, of when an assertion, when an appeal, when an explanation - all the whirl of organism Wittgenstein calls a "form of life".[35]

This shared form of life takes the form of shared basic, unreflective judgments - the same basic, unreflective grasp of how to apply concepts. Wittgenstein voices the worry that this would 'abolish logic'. Presumably, this is because it threatens the idea that logic treats of rational relationships between concepts, and the proposal under consideration would seem to reduce the question of which logical inferences are valid to a question of anthropology. But Wittgenstein makes the important distinction between 'judgments', which are basic responses that must be shared for any language to exist, and 'opinions', which are something for which adequate justifications can reasonably be asked.[36] While a system of

---

[35]Cavell (1969), p.52.

[36]See McDowell (1984), p.340, for this way of glossing the distinction.

measuring things depends upon at least some constancy in its results, one cannot actually justify a measurement by appealing to this constancy. One can refer to this constancy when one is describing the measurement system, but the process of obtaining and giving results will be governed by other norms that require proper justification to be given for any measurement. I shall say some more about this issue in the next section.

Practice or custom - consisting of regular, rule-governed behaviour resting on a foundation of shared basic responses learnt through the training that initiates one into a practice or custom - is the key to the possibility of understanding without interpretation. But it is crucial to remember that this suggestion is not being offered by Wittgenstein as a constructive answer to the question, 'How is meaning possible?'. Rather, it is meant as the second part of the process of correcting the misconception that makes that question seem pressing. First, the tempting, but ultimately doomed, attempt to understand our rule-following practices in a metaphysical realist (or when that fails, metaphysical anti-realist) way is shown up for what it is, then we are reminded of the features of ordinary rational, norm-governed practices.

7. Objectivity Secured

I have said that Wittgenstein can be seen as advocating a form of 'common sense realism'. I hope that I have successfully given some reason to think that the position is 'common sense' because of its rejection of the idea that it is compulsory that rule-following or meaning be warranted in either a metaphysical realist or anti-realist way. However, it might be less clear why this position deserves the label 'realist', and it is my aim in this section to show how the position does deserve this label.

A rejection of metaphysical realism (whether of a platonic or reductionist bent) is not the same as a rejection of our common sense view about the reality of rules or meaning. Wright and Kripke take it that the only way in which to explain the ratification-independent aspect of our intuitive idea of meaning is by appealing to platonistic imagery. But Wittgenstein's aim can be seen as one of separating the idea of objectivity or ratification-independence that is central to our conception of meaning from the fatally flawed mythology of platonism. Thus at *PI* 195 he writes:

> "But I don't mean that what I do now (in grasping a sense) determines the future use *causally* and as a matter of experience, but that in a *queer* way, the use itself is in some sense present." - But of course it is, 'in *some* sense'! Really the only thing wrong with what you say is the expression "in a queer way". The rest is all right; and the sentence only seems queer when one imagines a different language-game for it from the one in which we actually use it.

Wittgenstein's platonist interlocutor demands that grasping a sense must determine in a contractual (and not simply causal) way the future use of a word. The demand for this determination to take place in a 'queer way' is hopelessly mysterious. But Wittgenstein only disputes the platonist's demand that this determination take place in a 'queer' way, he does not dispute the demand that it take place at all. Because of the background of sensitivities and basic inclinations to action that underpin our meaningful use of words, there is a perfectly ordinary sense in which my grasp of a sense can be seen as containing my future uses of a word - in order to grasp the sense I must have these sensitivities and inclinations to respond in a basic way, and these will be what ensure that I apply the word correctly in future contexts.

There are two aspects of the normativity of meaning to which a critic might regard this interpretation of Wittgenstein as failing to do justice: its psychological compulsion and its objectivity. We will tackle each of these in turn. The platonist will claim to have an

account of the fact that we seem to be able to grasp the meaning of an expression in a flash and with utter certainty - it is simply a matter of one latching on to one of the ethereal rails that carry one on relentlessly to all future uses (*PI* 218-9). But this imagery is hopelessly mythological. How is the common sense realism about rules that I have been ascribing to Wittgenstein to do any better? It should come as no surprise that the answer comes from the fact that the practice of rule-following depends on basic, unreflective action in accordance with the rule. The psychological compulsion to apply a properly understood concept in a certain way, comes not from the intellect being engaged in some mysterious, but unfailingly compelling way, but from our acting in an unreflective way - "I obey the rule *blindly*" (*PI* 219). The lack of choice that we experience has nothing to do with compelling rails guiding us through conceptual space, but is a matter of our having been trained in instinctual action that precludes such choice. Because of the fundamental role of this action in providing the framework within which reason-based choices are made, choice is excluded when we respond to a rule in this basic way. It is even possible for Wittgenstein to accept that platonists are onto something when they introduce their imagery of guidance by ethereal rails. Their mistake is to think that these images are anything more than just images. So it would seem that Wittgenstein can meet the demand that there is a strong psychological compulsion correctly to deploy a concept that has been properly understood.

The second aspect of the normativity of meaning is its strong form of objectivity, which Wittgenstein highlights at *PI* 437:

> A wish seems already to know what will or would satisfy it; a proposition, a thought, what makes it true - even when that thing is not there at all! Whence this determining of what is not there? This despotic demand? ("The hardness of the logical must.")

Here Wittgenstein connects the phenomena of the intentionality of thought and expectation and the 'hardness' of logical necessity. In each of these cases, there is a 'despotic demand' to apply a concept in a constrained way - only certain things will satisfy the wish, only certain things will make a proposition or a thought true, and these are already determined. Now, platonists will, of course, claim that, in each of these cases, the initial deployment of a concept is connected in some mysterious way to the objects that fall under it. But can Wittgenstein, who rejects platonism, meet the challenge of accounting for this objectivity? The suspicion will be that his position cannot, because the role that is played by customs makes it a form of anti-realist constructivism.

We can, perhaps, best approach the idea of 'the hardness of the logical must' by considering mathematical and logical examples, where the despotic demand has its most natural home.[37] In such cases, it must be inconceivable to think of matters other than as they are. Thus we want to be able to say that the sum of five and seven would be twelve even if, in some counterfactual situation, we found it natural to give the answer as thirteen.[38] If an account of logical necessity fails to satisfy this demand, then it would seem to involve a considerable revision of our concept of logical or necessary truth - it would involve an abandonment of the idea that such truths are independent of our investigation or ratification.

Again the platonist might seem to have a simple answer to this demand. The sum of five and seven would continue to be twelve, even if we failed to be sensitive to it and gave the answer to the sum as thirteen, because the platonic rule-as-rail would still exist

---

[37]Mathematical and logical cases bring out the issue most clearly, but the same issue arises with any other concept. We could almost equally well run the argument that square objects would still be square, even if we found it natural to classify them with irregular quadrilaterals.

[38]See Cassam (1986) for this demand on an adequate theory of logical necessity.

in conceptual space. For Wright's anti-realist or Kripke's irrealist there is no fact of the matter about how the concepts involved should be deployed. Thus, if the counterfactual situation obtained and the communal assent conditions were such that thirteen was endorsed as the sum of five and seven, then thirteen would be the correct answer.[39] If Wittgenstein rejects the platonist's answer, can he avoid the constructivism that might seem to be the only alternative?

Wittgenstein's rejection of platonism is not a rejection of the ratification-independence of meaning and mathematical or logical truth. The answer to the challenge is that there is no conceivable way of formulating the counterfactual situation so that the sum comes out as thirteen. Having spurned the incoherent ambitions of platonism, the standards that determine the correct sum of seven and five are the standards of ordinary mathematical practice, and these standards determine the sum to be twelve. If the words 'seven', 'five', 'plus' and 'equals' all have their correct meaning, then it is simply false to say that seven plus five equals thirteen, regardless of what we might, or might not find, it natural to say. The only way in which 'seven plus five equals thirteen' could be taken to express a truth would be if the meanings of the words were different from what they are, but this would simply render the counterfactual situation undescribable and unevaluable. And this would be of no help to constructivists.

There might be a suspicion that this still leaves 'the hardness of the logical must' resting rather precariously on our particular practices. Thus, even if we cannot coherently formulate a counterfactual situation that challenges the necessity of logical truth, we still have the sense that what Wittgenstein's common sense realism offers by way of

---

[39]Although, see Craig (1975) and (1984) for discussion of the coherence of this position.

ratification-independence or objectivity is too insubstantial. This kind of objection might have some force if Wittgenstein was seen as giving a positive, constructive philosophical theory, but I have argued that he is not. His aim is to reject a metaphysical or philosophical realism because such a realist offers us an overly ambitious theory of the objectivity of meaning, which then disappoints and brings into question the whole idea that meaning can be objective. In the context of this debate, common sense realism rejects the idea that our following rules or meaningfully deploying concepts needs the sort of vindication that platonism and reductionism try to offer.

Although Wittgenstein rejects the idea that we can intelligibly see our concepts as matching up to a stock of utterly ratification-independent ethereal rules-as-rails, this does not mean that our deployment of concepts is not ratification-independent. At Z 358 Wittgenstein remarks:

> Then is there something arbitrary about this system? Yes and no. It is akin to what is arbitrary and to what is non-arbitrary.

Wittgenstein can be seen here to be making the point that a system of numbers or colours is arbitrary, in the sense that it does not reflect some underlying platonic framework of rules-as-rails. However, once that system, rather than some other, is adopted what one is required to do in deploying its concepts is perfectly non-arbitrary. It is not part of Wittgenstein's account that the application of rules and concepts is constantly open to determination by our inclinations, as constructivism would suggest. Basic action in response to a rule is itself essentially, and irreducibly, normative, and presupposes a conception of ratification-independent standards of correctness. Again, any circularity that might seem to be involved here is innocuous, because it is not part of a constructive account of the objectivity of meaning.

In order to forestall a worry that Wittgenstein's views import a wider idealism, we can also see how this point applies beyond the cases of mathematics and logic, for we can make a similar point in relation to talk about, say, physical objects. Part of the context in which we are trained to use words and deploy concepts is the material objects that surround us. These objects are intimately tangled up in the process of our acquiring the ability to use concepts concerned with physical objects. Once the over-ambition of metaphysical realism or platonism about concepts is rejected, we can see that, when asked what grounds we have for applying the word 'tree' to trees, it is perfectly acceptable to say that we apply it to such things because they are trees. The practice of using the word 'tree' has been born out of contact with examples of trees and rests on our similarity responses to those objects. This practice does not need a philosophical vindication in the form of some platonic ideal of a tree or some other metaphysical realist account of true treehood.[40] If it were objected that this account simply presupposes the existence of material objects and is therefore vulnerable to sceptical challenge, the response will be that Wittgenstein rejects the need to give an answer to the sceptic. In his view, the sceptic and metaphysical realist both share a misconceived view about what is needed to ground our claims to justified knowledge. We cannot coherently adopt the perspective from which the metaphysical realist wants to validate our knowledge claims and from which the sceptic wants to invalidate them, thus neither a sceptical challenge, nor an answer to it, can have any sense.[41]

In summary, then, we can make out the case that Wittgenstein's views on rules and the objectivity of meaning seem to satisfy the intuitive demands for psychological

---

[40]See Pears (1988), Ch.18, for more on this idea.

[41]See *OC*, passim, for Wittgenstein's discussion of these issues.

compulsion and a genuine, ratification-independent form of objectivity. The key to this is a recognition of the dialectical situation and of the fact that Wittgenstein is not offering a constructive theory of rule-following or the meaningful use of concepts.

8. The Charge of Quietism

At this point it would be worthwhile to tackle a common misunderstanding of the kind of interpretation of Wittgenstein's views on rule-following and meaning that I am offering, which leads to the charge that Wittgenstein is guilty of the philosophical crime of quietism. On this view, Wittgenstein is seen as unearthing, through his rule-following considerations, an important set of philosophical problems that stand in need of a constructive exercise in theory-building. Depending on the precise details of the interpretation of Wittgenstein's views that is in question, he can be seen as guilty of one or other of two related failings.

On the first reading, Wittgenstein recognises the problems that he has brought to light as genuine ones, but then refuses to answer them or, worse still, suggests that it is impossible to answer them.[42] Gareth Evans has made this kind of objection against Wittgensteinians:

> [B]ut it is surely equally deplorable for students of those ideas to act as intellectual Luddites, dismissing the entirety of a sophisticated and developing tradition without a detailed consideration of its findings, and an alternative account of the enterprise to which the obviously compelling distinctions and observations it contains do properly belong.[43]

Simon Blackburn characterizes quietism as "the attitude which ... urges that at some particular point the debate is not a real one, and that we are only offered, for instance,

---

[42]Wright (1992), Ch.6, and Boghossian (1989), pp.507-8, are examples of those who have read Wittgenstein in this way.

[43]Evans (1981), p.326.

metaphors and images from which we can profit as we please".[44] As James Conant has pointed out, this sort of definition is importantly unclear.[45] It can be developed in two possible directions. On one interpretation, which Conant dubs 'dismissive neutralism', quietism accepts that there is a distinctive philosophical standpoint from which to raise questions, but urges us to renounce it. In this context, the charge of quietism against Wittgenstein would amount to the claim that he recognises that there is such a thing as the task of constructive philosophising, but renounces doing it. If this was Wittgenstein's view, it would be easy to sympathise with Evans's frustration with its evasive and Luddite quality. Equally, it would seem to be open to us not to follow Wittgenstein in his renunciation, but instead engage in the perfectly intelligible task of philosophical construction. This would have the consequence of rendering the account of objectivity that I have sketched at best inadequate and, at worst, downright unstable.

However, there is another way of glossing the idea of quietism. On this view, one should not leave off philosophical debates once they reach a certain point; rather one should not get started on them at all, because the whole debate rests on a serious misconception. Hilary Putnam can be seen as attributing this view to Wittgenstein:

> [H]e did not wish to sweep problems under the rug; what he was rather trying to do is see just what picture 'holds us captive' - to find the roots of our conviction that we have a genuine problem, and to enable us to see that when we try to state clearly what it is, it turns out to be a nonsense problem.[46]

If this does amount to a form of quietism, it clearly is not the form of renunciationist Luddism that has rightly earned a bad name. This approach requires a considerable

---

[44]Blackburn (1984b), p.146.

[45]Conant (1996-7), p.210.

[46]Putnam (1996), p.252.

philosophical effort to understand our cognitive practices and the questions (or pseudo-questions) that we ask ourselves. It should also be clear that this position does not advocate as a strategy an evasion of problems. Rather, the problem itself is tackled head-on and resolved, not, admittedly, by constructive theory-building, but by showing the problem to rest on a mistaken conception. On this view, it will be an important part of our attempts to understand our cognitive practices that we examine the problems that we face, in order to see whether they are really the products of a deep misconception.

I said that there were two ways in which a philosophical position might end up facing the charge of quietism. The second is by trying to effect a different response from the one that we have pursued to the unhappy situation that we had arrived at at the end of Section 5. One might take it that the failure of reductionist theories of meaning and the incoherence of meaning irrealism suggest that the required theory must be realist (because of the incoherence of the anti-realist or irrealist option) and must be non-reductionist (because of the failure of reductionist accounts).[47] The difficulty for such an approach, once one has taken on the burden of constructive philosophizing in a realist mode, is that of steering a course between the twin dangers of reductionism and quietism. We have a fairly clear idea of what a reductionist attempt at theory-building looks like (even if, ultimately, it fails), but we have a less clear view of what a non-reductionist project might look like. This, in itself, is not a compelling problematic, as we might simply be relatively unadvanced in our understanding of these matters. The real problem comes when such an approach tries to maintain its distance from quietism. Non-reductionists will want to say

---

[47]This is the conclusion that Boghossian (1989), especially section VI, draws from a discussion of the state of the post-Kripkean debate. He does not address the sort of approach that I have advocated, perhaps because he sees the terms of the debate as set by Kripke's formulation of the issue and so shares the latter's blindness to this option.

that at some point no further analysis is possible. If they try to take Wittgenstein as their guide in developing a non-reductionist form of realism, that point might come pretty quickly. But this now lays them open to the charge of quietism - they have accepted the challenge of answering philosophical problems, but then call a halt to the project at some apparently arbitrary point.

In contrast, on my interpretation of Wittgenstein, we can give a reasoned justification for why one should not give constructive answers to a philosophical problem, namely by pointing out that the problem itself may rest on a misconception that needs to be corrected. The non-reductionist realist cannot give this response, but is left embarrassingly open to the charge of illicit quietism. We can, perhaps, find an example of a philosopher ending up in this predicament in John McDowell's paper "Wittgenstein on Following a Rule".[48] McDowell seems to waiver between a position like the one that I have adopted and the temptation to give a non-constitutive account of the role of community in following a rule. But because of his desire to avoid a substantive account of the role of community, it is unclear what role he thinks that the community actually does have.[49]

In the light of this second route to quietism - via a failed non-reductive realism - it should again be stressed that my discussion of the role of practice in following a rule is not meant to suggest any kind of constructive theory of meaning. Rather, it is meant to offer 'reminders' of what our normative practices actually rest upon; reminders that are

---

[48]McDowell (1984).

[49]See Boghossian (1989), p.544, for this charge. McDowell (1992), p.49, explicitly rejects the idea that Wittgenstein should be read as giving a constructive answer to the question 'How is meaning possible?', but admits in a footnote that his exposition in the earlier paper does lay him open to the accusation that I have levelled at him.

meant to dissipate some of the feelings of vertigo that might be induced by the realization that our practices are not grounded on platonic rules-as-rails.

9. Is There Anything Left to Do?

Some philosophers, inspired by the lessons that they find in Wittgenstein's work, go on to embrace any quietism that they find there. In the light of this they will see the role of philosophy as reduced merely to describing the differences between 'language games'.[50] It might, thus, be thought that an attempt to engage in anything like meta-ethical theorizing will be deeply mistaken, and doubly so if the attempt is launched from Wittgensteinian foundations.

This lack of ambition seems entirely unwarranted. The sort of realism that I have found in Wittgenstein involves a rejection of the fabrication of philosophical theories to underpin our rational practices. This certainly precludes attempts to philosophize in the way that anti-realist constructivists, platonic realists or reductive realists do, but it does not rule out all philosophizing.[51] Nor are we left simply with the task of trying to unearth misconceptions that hold us in their sway and drag us into fruitless metaphysical debates. There is also the task of examining our rational practices from within and trying to find out what counts as a good reason for thinking, saying, believing or doing something. This task can be entirely detached from the distorting perspective of those views that seek to give these activities an external grounding.

---

[50]See, for example, Johnson (1989).

[51]There is even scope for theorizing about meaning, as is demonstrated by attempts to develop 'modest' theories of meaning that do not try to leave behind all semantic or normative concepts. Examples of such modest theories are McDowell (1977), Wiggins (1980) ,various of the essays in Davidson (1984), and Davidson (1990). Modesty is defended in McDowell (1981b) and (1987).

There is also scope left for asking questions about the objectivity of the judgments within a particular discourse, provided that this is not seen as a matter of examining whether the concepts deployed in the discourse match up to some platonist ideal or can be reduced to something naturalistically respectable. We can, for example, ask whether the ethical domain is more like natural science or mathematics in the discipline that it exerts over agents' deliberations, or more like some domain that we do not think of as particularly objective - or objective at all - such as discourse about humour or gustatory taste. Before this task can be effected we need to set up some kind of framework within which this process of assessment can take place. But this must be a framework that is free of the ethical anti-realist's assumption that the proper standard of objectivity must be one that takes the natural sciences as its model. One of the lessons of this chapter must be that when we think this position through we are actually forced to the absurd view that even the meanings of our concepts are unreal, a conclusion that seems to lead to a global idealism.

We can also deploy another argument against the Wittgensteinian who wants to eschew philosophical theorizing in ethics. This response might concede, purely for the sake of argument, that Wittgenstein had a legitimate hostility to philosophical theorizing and thought that what one had to do was to try to diagnose and dissolve philosophical problems. But we would then point out that Wittgenstein's conception of philosophy is unusually narrow, since, in his view, philosophy is really only concerned with questions about logic and language. There are many other topics that lie outside these areas that are scrutinized in philosophy departments, and ethics, aesthetics, political philosophy and scientific methodology are examples of these. These are subject matters that are highly theoretically complex and that demand reflective investigation in order for us to get a

clearer idea of the concepts that are involved. In these areas there are pressing theoretical questions about what constitutes good argument - or, indeed, questions of whether there are such things as good arguments in these domains at all - that are not the inventions of the diseased minds of philosophers, but are the products of our everyday use of these concepts. We can, if we like, stick with the narrow definition of philosophy that, on this interpretation of his views, Wittgenstein is seen as offering, and then simply drop the claim that what we are doing, when we deal with these theoretical questions, is moral *philosophy*. Instead, we could see ourselves as engaged in theoretical reflection on our ethical concepts.

If our Wittgensteinian opponents still quibble at this, then they seem to be attributing to Wittgenstein a fairly extreme form of anti-intellectualism and obscurantism.[52] But this attribution would seem to be hard to sustain. Hilary Putnam cites Stanley Cavell as raising the point that, when Wittgenstein uses the notion of 'ordinary language', he does not mean a language that is simple or everyday, in a way that scientific or other specialized language is not.[53] Rather, the ordinary (which can include the specialized) is to be contrasted with the philosophical, and Wittgenstein had no intention of interfering with the ordinary. So, if one feels more comfortable eschewing the philosophical, then we could declare the philosophy in this thesis over; we now move on to discuss what standards there might be for good argument in the ordinary discourse of ethics. For my part, I am happy to think that this activity counts as philosophy of a non-pernicious sort.

---

[52]Winch (1991) seems to be in some danger of falling into this trap in his discussion of political obligation.

[53]Putnam (1996), p.260.

10. Conclusions

In this chapter, we have entered some way into the wider debate about realism to see if there is a coherent way of thinking about objectivity and reality that might offer an alternative to the conception of reality to which ethical anti-objectivists and non-cognitivists seem to adhere. By looking at Wittgenstein's discussion of following a rule we have found what I hope is a compelling set of considerations that challenge the coherence of the demand that the objectivity of meaning must be validated by something external to our normative practices. With the legitimacy of this demand undermined, we can see that the objective correctness of judgments does not rest on their matching some platonic ideal or their being the result of some reductively defined ideal disposition. Instead, judgments are objectively correct if they are based on reasons that, in an ordinary sense, count as good reasons. It should be obvious that this conception of objectivity does not immediately suggest that there will be a problem for the objectivity of ethics, in the way that the ethical non-cognitivist's conception of objective reality did. Human sensitivities and concerns are seen as underlying the practice of using words and concepts in meaningful ways. This also suggests that the non-cognitivist moves too swiftly in concluding that, since beliefs are simply passive reflections of a motivationally inert reality, putative ethical beliefs cannot have the appropriate link to action.

These Wittgensteinian considerations would seem, then, to undermine some of the knock-out arguments that non-cognitivists use against ethical objectivism and cognitivism. But they might not seem to get us very far towards establishing that ethical objectivism is, in fact, the correct position to adopt. We might, with some justice, think that it is not enough simply to hold that if a discourse has an assertoric form, then it is perfectly in

order to regard it as realistic or objective.[54] There still seems to be plenty of scope for the project of examining domains to see what the appropriate standards of good reasons are within them. Wittgenstein might have baulked at the idea that this could involve an assessment as to how objective the judgments within a given domain are. However, it does seem that one thing we might want to do is to see whether a target discourse (in our case, ethics) is more similar to physical science or mathematics in the rigour of its standards of correctness, or more similar to, say, gastronomy or humour. In the next chapter, I shall try to set out how one might develop a framework in which to assess the level of objectivity exhibited by judgments within a domain. This framework will have to be compatible with the broad constraints that we have unearthed in this chapter. We will try to go beyond a crude, and rather caricatured, Wittgensteinian position that regards all cognitive practices as equally objective if their surface syntax is of an assertoric form. But we will also want to avoid the anti-realist assumption that there are only really two categories of discourse - the genuinely realistic and the rest, which is to be treated in expressivist or projectivist terms. We will also want to conduct this enquiry in a way that is not question-begging either in favour or against objectivism with regard to a particular domain; for example, we will want to avoid a view, such as Crispin Wright's, which regards an anti-realist treatment as the default position.

---

[54]We will see in Chapter 5 that an influential attempt to move from Wittgensteinian considerations about objectivity to moral realism does not demand much more than this.

Chapter 4

A FRAMEWORK FOR ETHICAL OBJECTIVISM

1. The Limited Lessons So Far

In the last chapter, we were concerned with broad issues about realism and objectivity in general. The purpose of this chapter is to return us to the discussion of these issues as they relate to ethics. In particular, my aim will be to refine the terms of the debate, so that we can go on to examine two types of ethical objectivism that both start from an acceptance of a broadly Wittgensteinian conception of objectivity: the virtue-centred moral realism of John McDowell and my own objectivist account of ethical judgment, which draws on Kantian ideas about the role of norms of reason. Before we can look at these theories in detail, we need to draw out more fully the implications of our discussion of Wittgenstein for the debate about ethical objectivism. For this reason, this chapter will not offer much in the way of substantive argument, but will set the scene for the rest of the thesis.

In Chapter 2.5, I identified three key groups of problems that we can see ethical anti-objectivists or non-cognitivists raising for ethical objectivism and cognitivism. These problems were: the prescriptivity of ethical judgment and its incompatibility with ethical judgment being a form of belief; the lack of appropriately produced rational convergence in ethical judgment; and the incompatibility of the essential contestability of ethical concepts with objectivity.[1] I argued that the third of these problems was derivative from the second. So, this leaves us with problems surrounding 'queerness' and the prescriptivity

[1]And these problems also arise for ethical objectivism that is not of a cognitivist variety.

of ethical judgment and with issues concerned with convergence in ethical judgment. What light does our reconsideration of a generalized conception of objectivity shed on these two issues?

The ethical anti-objectivist was seen to hold a conception of objective reality that is hostile to the possibility that ethical judgments might exhibit objectivity, or, slightly more ambitiously, the possibility that they might be capable of being true. This conception of reality is such that a judgment counts as objectively correct only in so far as it reflects features of a motivationally inert world that is entirely independent of our sensitivities and interests. It is this conception that gives rise to the charge that ethical judgments cannot be objectively correct, because they are clearly not judgments about features of the world conceived of in this way. Because the world is conceived of as utterly independent of our interests, any objectively correct judgment about it will have to take the form of a motivationally neutral belief. Yet ethical judgment seems to have an intrinsic link to action, and so it is impossible for an ethical judgment to be an objectively warranted belief. This, then, is the problem with the prescriptivity of ethical judgments.

The second sort of problem concerned the supposed inability of ethical judgments to match up to appropriate standards of rational convergence in judgment. The favourable case is of judgments about the world as described in perspective-free physical science. Here we converge in our judgment that p, and this convergence is explained by the fact that p pertains independently of our judgments and causes us to come to judge that p. The challenge to ethical objectivism is that where there is convergence in ethical judgment it is explained, not by our judgments responding to features of the world, but merely by our sharing certain attitudinal capacities. The ethical features to which we might claim to

respond are simply the shadows cast by our own responses, and thus do not have the necessary independence to qualify as things to which we are genuinely responsive.

This, then, is a very rough and ready reminder of the two key challenges to ethical objectivism. We can already see that our discussion of Wittgenstein suggests ways in which the ethical cognitivist or objectivist might start constructing replies to these challenges. In both cases, the response will start from the thought that Wittgenstein shows that the strongly metaphysical, or philosophical, realist standard of objectivity, which ethical non-cognitivists and anti-objectivists wield, is suspect. Our ability to apply concepts in an objectively correct manner is not warranted in the way that this conception demands. We do not use concepts in an objectively meaningful way because we are responding to some hyper-objective platonic idea, nor because we are exercising a disposition that can be reductively defined in strictly non-normative terms. Rather, the demand for this kind of validation of our use of concepts is to be rejected as incoherently over-ambitious. Our ability to use concepts in a meaningful way is a result of our being inculcated into an irreducibly norm-governed concept-deploying practice. What implications does this understanding of objectivity have for the prospects of ethical objectivism?

2. Wittgensteinian Realism and the Anti-Objectivist Challenge: 'Queerness'

If we look first at the problem of the prescriptive nature of moral judgments, we can start to see the shape of a possible objectivist response to the anti-objectivist challenge.[2] The

---

[2]This is really a very rough sketch to serve merely as an advertisement for what is to be discussed more fully in Chapters 5 and 6.

assumption that beliefs must simply have the 'direction of fit' of matching a neutrally characterized world starts to look less secure once we have in mind our new understanding of the objectively warranted use of concepts.[3] On the Wittgensteinian account that I have been exploring, what justifies a belief as objectively warranted or true is not simply its passively reflecting the world, but my being able to deploy the concepts involved in a competent way when required to do so. This ability will not be explained simply as a passive reflection of the world, but as possession of a complex set of recognitional, cognitive and behavioural abilities. It will be dependent on my participating in a practice of using the relevant concepts, and this participation will be the result of my acquiring certain sensitivities and understandings of salience. It is natural to think that this process will be holistic in nature and will also involve the ability to reason with the concepts that I have acquired. Such reasoning will require that my beliefs do not just have a passive, world-reflecting direction of fit, but may also have a component that has a direction of fit that requires the world to change - the direction of fit traditionally associated, not with beliefs, but with desires.

This point comes out most clearly in a case where the acquisition of a belief has more obvious normative implications, for example, the case of a mathematical belief. If I come to believe that two plus two equals four, my belief will be a full and genuine one only if I am also prepared (should it prove necessary) to adjust my other beliefs to bring them into line with the implications of this new belief. Although this comes out most clearly in a mathematical case, other types of belief might also seem to have this

---

[3]See, for example, Smith (1994a), pp.111-9, for the Humean view of how desires are states whose 'direction of fit' is to get the world to fit the desire, and beliefs are states whose 'direction of fit' is to fit the world (although, unlike most Humeans, Smith goes on to develop a form of moral cognitivism).

characteristic. Even my humble belief that there is a cat on the mat requires me to change the other beliefs that I have, for example that there are no non-human mammals in the room or that the mat is simply a hologram and not a real, cat-supporting mat. In these cases, it would be crazy to suppose that I only have a reason to change my other beliefs if I have some desire, for example, to have coherent beliefs.

Given the Wittgensteinian framework, it is a relatively small step to suppose that the acquisition of a belief need not only give us reason to change other propositional attitudes that we have, but might also give us reason to act in certain ways. The examples that Wittgenstein uses of people participating in norm-governed activity often involve people engaged in action and not simply mental cogitation. The natural responses on which our norm-governed practices depend can include non-verbal responses. Thus we can start to see how it is possible, on a Wittgensteinian approach, to say that a moral judgment could be intrinsically connected to a reason to act, without its involving a desire - to have a certain belief just is to have a reason to act in a certain way. One might simply point out that on Wittgenstein's account the basic metaphysics simply does not presuppose a picture of inert reality (reflected passively by belief) and 'ert' (as one might put it) appetitive states of subjects (expressed in desires). This is an incredibly loose sketch of how one might start to respond to the challenge of the prescriptivity of moral judgment. Clearly, a very great deal more will need to be said before this looks anything like a proper account of the way that moral judgments might give one a (motivating) reason to act without involving a desire, or other non-objective or non-cognitive ingredient. We will see how this approach might be developed in the latter part of Chapter 5 and in Chapter 6.[4]

---

[4]And this sort of argument will not only be useful for ethical cognitivism, but also for ethical objectivism more broadly. The basically Humean idea that desires are required for reasons to act

3. Wittgensteinian Realism and the Anti-Objectivist Challenge: Convergence in Judgment

The second major ethical anti-objectivist or non-cognitivist challenge concerned the issue of whether ethical judgments could exhibit a form of objectivity-sustaining convergence in judgment. How does this challenge stand up in the light of our discussion of Wittgenstein?

One might take the common sense realism that I have attributed to Wittgenstein as offering a very obvious payoff in those areas where we use apparently realist, truth-invoking methods of assessing judgments but do not seem to be talking about bits of the fabric of the world (nor anything reducible to such bits). Those areas are, for example, ethics, mathematics, and semantics. The anti-objectivist wants to say that in these areas, since there are no metaphysically respectable objects to which we are referring, we must simply be projecting our own attitudes onto the world. However, we have seen that this notion of 'metaphysical respectability' is suspect. The lack of properties identifiable by physical science to ground our ethical judgments can be seen as presenting as little problem for objectivism or cognitivism about ethics as the lack of platonic mathematical objects does for realism about mathematics. But is it really that easy? Can we just say that, because ethical discourse involves the attribution of truth to ethical judgments, objectivism, or even realism, about ethics is warranted?[5]

Sadly, I suspect that the answer to this question is 'No'. To see why, we have only to think of those areas where we also use seemingly truth-involving talk, but are much less willing to draw the inference that the subject matter fully warrants objectivism. Examples

---

can be equally problematic for forms of ethical objectivism that are not cognitivist or realist.

[5]See Putnam (1996), p.264, for this suggestion.

of such areas can come in varieties that support an error theory for the domain in question, or that support a projectivism or quasi-realism. Examples of the former would be talk about the magical, the astrological and (some would argue) the divine. Examples of the latter would be the humorous or the tasty. It does not seem to presuppose a flawed conception of objectivity to ask whether there is some difference in objectivity between ethics and discussions of the humorous or tasty. Indeed, it seems only too natural a question to ask. If there is no significant difference between the level of rational discipline exerted by ethics and that exerted by talk of the humorous or (worse still) the tasty, then ethical anti-objectivism will have won through without relying on an untenable metaphysics. So, there is an urgent need for a framework within which we can start to address the question of how to distinguish subject matters according to their objectivity, even if they all exhibit an assertoric form. This will need to be done in a way that does not presuppose any dubious metaphysics. I shall spend the rest of the chapter trying to construct such a framework.

4. 'Response-Dependence' and Objectivity

Ethical judgments fall into the category of judgments that exhibit an assertoric form and apparently aspire to objectivity, but cannot sensibly be construed as trying to respond to bits of the world as understood by natural science. One way of trying to account for this is to claim that such concepts are 'response-dependent', and that objectivity is compatible with such response-dependence.[6] Concepts are response-dependent if an explanation of

---

[6]Examples of this idea being deployed in the case of ethics can be found in Lewis (1989), Smith (1989), Johnston (1989), Smith (1994a), Putnam (1981), Ch.9, Wiggins (1998), Chs. III-V, and McDowell (1979), (1981a) and (1985b).

them must make essential appeal to the attitudes or responses of subjects who use them. Such an approach will stand opposed to a projectivism or error theory because it will claim that, despite the role of subjects' responses, the use of the concepts is sufficiently disciplined for talk of objective correctness, or of truth and falsehood, to be in good order.

There seems to be a natural affinity between this sort of response-dependent objectivism and a common sense realism. Both reject a tendency towards the hyper-objectification of the conditions for the correct use of a concept, and both accept the key role of subjects' sensitivities in accounting for the objectively correct use of a concept. However, the two positions are probably best not regarded as identical, since response-dependent objectivism will typically consider the reactions of subjects to be essentially involved in a specification of a response-dependent concept.

The difference can be illustrated in the following way. A common sense realist will say that the reason that I am justified in saying that there is a tree in front of me is that there is a tree in front of me, even though my ability to apply the concept of a tree is dependent on my belonging to a practice of using the concept 'tree'. If I failed to accept this justification, it would not count as a legitimate form of further justification to point to my membership of the practice. Instead, it would have to be pointed out that (for example) the light is good, I have not been drugged and so on, and thus I have no good grounds for doubting that there is a tree present before me. On the other hand, someone espousing response-dependent objectivism about ethical value might suggest a definition of value such as the following, which Mark Johnston has suggested:

> X is a value iff practical reason is on the side of valuing x, i.e., the deliverances of good practical reasoning support the conclusion that x is a value.[7]

Here, someone's failure to recognise a value might appropriately be met by pointing to the right-hand side of this biconditional. Although a specification of the sense of a response-dependent concept involves an appeal to agents' responses, the reference of such a concept does not. Thus, like common sense realism, the position aspires to avoid falling into a form of empirical idealism, according to which the subject matter depends for its very existence on responses as they actually are. The response-dependent framework seems to be a natural way to structure an investigation of where to locate ethics on the scale of objectivity, given the overall conception of objectivity that I have advanced.

The key question facing a response-dependent objectivist theory is what the relevant responses are that determine the correct application of the concepts in question. Crudely, the answers to this question will come in two varieties: reductionist and non-reductionist. In the case of ethics, reductionists will try to pick out the relevant responses in non-normative, non-ethical terms, for example by equating what is valuable with what would be chosen by a fully informed and rational Ideal Observer. Such theories seem to me to be deeply flawed. Briefly stated, there are two main groups of reasons for this.[8]

---

[7]Johnston (1989), p.154. I only introduce this as an example and do not mean to endorse the proposal at this stage.

[8]McGinn (1983), pp.145-54, and Blackburn (1985), offer a number of objections to attempts to explain the objectivity of ethical judgments on an analogy with secondary qualities, that is, roughly, in response-dependent terms. Some of these objections are aimed at the particular analogy with secondary qualities, such as colour. As I do not wish to pursue this analogy, I shall not deal with such objections. The rest of the objections seem to be at their most impressive when directed at reductive response-dependent theories. Wright (1988a), pp.5-11, offers plausible answers to these points. Wright, however, does go on to argue that nonetheless the analogy that the ethical realist really ends up with is with primary qualities, and hence they are open to the old problems raised by non-cognitivists and error theorists. I shall not consider this argument. In the end, from my point of view, it does not really matter if this argument works, because, on the common sense realist picture that I have sketched, it would not be so obviously problematic to regard ethical properties as primary qualities.

First, we can repeat the sort of objection that was raised against reductive dispositional accounts of meaning in Chapter 3.3, and here is Alan Gibbard raising just such an objection:

> Fix an account of what conditions for judgment are ideally favourable, and select a response. Call the conditions C and the response R. The analysis says that normative judgment J is correct if and only if a person in condition C would have response R. We can imagine someone, though, who admits that response R was evoked in conditions C, but who disagrees with judgment J and disagrees coherently. He makes no mistake of language or logic. Rather, he has a different account of what conditions for normative judgment are ideally favourable.[9]

Because the ideal conditions are non-normative and non-ethical, it is perfectly intelligible, and shows no disrespect to values, to ask "So what?", when told that a certain judgment is reached under ideal conditions.

The second set of objections concerns cases of disagreement and the treatment of them that a reductive response-dependent theory can offer. Colin McGinn offers such an objection when he writes:

> [T]here is no reaction-independent criterion for [colour] ascriptions. This point shows that moral error, on the dispositional theory, could consist at most in a failure of conformity of one person's moral reactions with the reactions of others: the theory cannot allow that a whole community might be in moral error, or that a solitary judger might make moral mistakes, since this would require some standard of correctness external to that provided by an essentially arbitrary norm of moral reaction.[10]

The objection is that, when correctness in the application of a moral concept is reductively defined in terms of the responses of one's moral community, there can be no coherent sense in which the community can go wrong or suffer a decline in its moral standards. This seems to be deeply implausible.[11] This leads on to a further kind of objection, namely that

---

[9]Gibbard (1990), p.186.

[10]McGinn (1983), p.151. See, also, a similar objection in Blackburn (1985).

[11]Again this sort of point should be familiar from our discussion, in Chapter 3.3 and 3.5, of communitarian theories of meaning. In fact, the use of rigidification of the community's actual responses can alleviate this problem (see Wiggins (1998), p.206, and Johnston (1989), pp.140-1,

on the reductionist picture it is impossible for there to be reasonable disagreement, and hence impossible to conceive of a society better than one's own. Again, this implication is of dubious plausibility. Simon Blackburn makes these points thus:

> I believe that John McDowell's views on virtue suggest that with increasing virtue comes an increasing approximation to the mathematical case, so that the virtuous man is eventually distinguished by a certain inability to see how reasonable men can differ. This, I must say, represents a value with which I find it hard to enthuse... with us who are less exalted there exists a lively sense of the objectivity of ethics alongside a lively awareness of alternative points of view... Second, one of the essential possibilities for a moral thinker is that of self-criticism, of the thought that our culture and way of life leads us to corrupted judgment... And it is not easy to see how this is to be construed if objectivity is somehow 'based on' consonance in a form of life.[12]

We have already seen how this kind of problem emerges in the context of a reductive dispositional account of meaning, and have, in part, rejected that theory for this very reason. Within the context of the broadly non-reductionist approach that I am suggesting, and given these problems with reductive response-dependent theories, it makes sense to move swiftly on to consider the non-reductive variety.

Non-reductive response-dependent theories do not try to provide a reduction of the ethical to anything non-ethical. This means that their definitional biconditionals will be, to a greater or less extent, circular, as is Mark Johnston's definition of value mentioned a few pages ago, which makes essential reference to the deliverances of practical reasoning. This need not be problematic, for the goal of such non-reductive accounts is not to provide an analysis of ethical concepts, but to elucidate these concepts and their connections with each other and with other concepts. Provided that the

---

and Wright (1988a), pp.8-9).

[12]Blackburn (1981), p.171. I think that Blackburn is mistaken in thinking that McDowell holds some kind of reductive response-dependent theory of value. This does not mean that McDowell has no problems in this area - in the next chapter I argue, in effect, that he faces a dilemma of either facing the objection that Blackburn suggests, or offering far too weak an account of objectivity.

explanatory circle is suitably wide, there is a good chance that this goal might be achieved. I want to adopt the broad approach to conceptualising response-dependent forms of ethical objectivism that we find in Johnston's proposal and in David Wiggins's moral cognitivism, which we touched on briefly in Chapter 2.7. Here the key idea is that the objective correctness, or truth, of ethical judgments can be explicated in terms of those judgments matching up to the demands of practical rationality.

On Wiggins's account, a subject matter is objective if it admits of answers that are substantially true:

> It is sufficient for some judgment that p to be substantially true that one could come to know that p. One can come to know that p only if one can come to believe that p precisely because p. And one comes to believe that p precisely because p only if the best full explanation of one's coming to believe that p requires the giver of the explanation to adduce in his explanation the fact that p.[13]

And this explanation must take the form of showing that "for this, that and the other reason (here the explainer specifies these) there is nothing else to think but that p".[14] So the question of whether a subject matter is objective is a matter of whether it admits of such vindicatory explanations in a sufficiency of cases. The utility of Wiggins's formulation of the issue is that the idea of 'there being nothing else to think' has the potential to be equally applicable to subject matters for which objectivism might be in order, but for which one might be wary of advocating full-blooded cognitivism.[15] The affinity between Wiggins's demand for explanations that show why there is nothing else to think and Johnston's definition of the concept of a value as what practical reasoning is on the side of should be clear.

---

[13]Wiggins (1990-1), p.66.

[14]Wiggins (1990-1), p.66.

[15]See Moore (1996) for a useful discussion of this requirement.

It should be obvious that what has been said so far offers only the start of a way of understanding ethical objectivity. In particular, we will need to have a very much better understanding of what the appropriate standards of practical rationality and rational criticism are that can leave one with nothing else to think. I shall pursue the task of developing that understanding over the next two chapters. However, before that we need to deal with the suspicion that this approach to objectivity is inadequate because it ignores the necessity of a further constraint on rational judgmental convergence, namely the requirement of width of cosmological role.

5. What Role for Width of Role?

Within the framework that I have suggested, the objective correctness of a judgment will, at least in part, be a matter of its surviving rational criticism that is, in some sense, internal to our rational practices. The key question is how 'internal' the standards of this criticism and of practical rationality are. Clearly, a Wittgensteinian common sense realism requires us to abandon the notion that these standards are given a grounding that is metaphysically over-ambitious. However, there still seems to be scope to demand some account of the standards that such criticism is to meet if it is to sustain objectivity in any full-blooded sense. Here we might seem to be drawn back to Crispin Wright's way of conceptualizing the discussion of objectivity, which demands that objective judgments be responses to things that exhibit wide cosmological role. The common sense realist insists that, if a discourse has the features of truth-aptness, then it should be treated in a realist way. But is this not oblivious to Wright's claim that the anti-realist need not deny truth-aptness to discourses that seem to be truth-apt, but should instead ask what kind of truth predicate

is in play?[16] Given the common sense realist picture, there still seems to be an intelligible question left unanswered about the level of discipline (or objectivity) exhibited by a domain. The issue that this section deals with is whether this question is best answered by adding a requirement of width of cosmological role, or indeed some other width of role test. I shall argue that width of role of whatever kind is, at best, a reassuring feature of a domain's objectivity and is neither a necessary nor a sufficient condition for objectivity.

I argued in Chapter 2.7 that the key issue for those proposing objectivism about a domain is that there be convergence in judgment and that the convergence be brought about in the right way. In Wright's view, such convergence in the judgment that p counts as objectivity-sustaining only if the convergence is explained by the fact that p, and the fact that p has a suitably wide cosmological role. Wright might seem to have struck upon a key intuition about objectivity. The thought is that, for all its truth-apt appearance, talk about, for example, the astrological fails to be objective because the properties referred to have no causal role, or indeed any other role, since they do not exist. The discourse of physical science gives rise to objective judgments about properties that we are confident exist because they have a wide causal role. The idea is that we can be confident of the independent existence of a property if it features, not just in the explanations of one discourse, but in others as well. And to achieve this it must have a genuine causal role.

However, from our point of view this requirement seems problematic for at least two reasons. The first is that it seems to be a requirement for realism rather than objectivity more broadly understood. I am not disputing that ethical judgments are unlikely to meet a standard of objectivity conceived of in strongly realist terms; the question is

---

[16]Indeed, Wright sees his project as fulfilling the Wittgensteinian aim of investigating the differences between 'language games' (Wright (1992), p.204).

whether they might meet a less ambitious, but nonetheless genuine, standard of objectivity. Granted that standards of reality are one species of standards of objectivity, there still remains a second problem for this proposal. In Chapter 2.7, I suggested a dilemma for this approach: either it is to be interpreted strictly, so that the class of causally efficacious properties is restricted to those picked out by physical science; or it is interpreted in a fairly liberal way, so that non-reductive monism would allow a wide range of properties to have a causal role. In the light of our discussion in the last chapter, it would appear that the first option presupposes the dubiously over-ambitious metaphysics of metaphysical realism.

If we consider the second horn of the dilemma, the question is, Might this not leave the width of role requirement too weak? My suspicion is that it does. If the test is relaxed in this way, then ethical properties might count as genuine because they give rise to ethical beliefs and judgments, which, in turn, give rise to actions. If a non-reductive or anomalous monism is true, then any account of what happens in the world would be incomplete, unless it referred to these ethical properties. So, there is a certain plausibility in suggesting that ethical properties do have a causal role that ramifies quite widely. The problem is that, as I argued in Chapter 2.7, a similar story might be told to show that magical or astrological properties are genuine, provided that there was sufficient convergence in judgment among 'observers' of the magical or astrological and these judgments led to action.

One response to this worry might be to stay with a causal-explanatory conception of width of role, but seek other ways in which the magical or astrological failed the test. For example, it could be pointed out that astrology tries to offer a predictive causal-explanatory theory that rivals natural science, but that it fails miserably in this endeavour.

Its predictions simply cannot be specified to anything like the degree of precision needed to render them testable. Astrology gets into difficulties because it tries to engage in the game of causal explanation, but is totally unable to match up to the norms of good reasoning within that discipline. But there seem to be two problems with this kind of response. The first is that one might suspect that there are dubiously objective domains that do not fall foul of this problem, because they do not try to take on natural science at its own game. Examples of such domains might be the humorous, the divine, and the ethical. These domains cannot be charged with failing to match up to the standards of good causal explanation, since they do not claim to be playing that game. The second, and more pressing problem, is that this sort of response does not actually seem to have much to do with the *width* of the cosmological or causal role. The real problem for astrology or magic is not the narrowness of any role that their putative properties might have, but their failure to offer explanations that are any good compared with those offered by competing, and far more successful, explanatory schemes.

Width of cosmological role does not, after all, seem to offer a satisfactory extra constraint on rational convergence in judgment. So, the question might be asked whether we can develop this intuition about width of role in a more fruitful way if we abandon the causal aspect of the requirement. That the properties that are picked out in a certain domain play a causal role is certainly one way that we can be confident that they are independent of our judgments about them. But there may be other ways in which this independence can be established that are more friendly to an objectivism more broadly characterized. Width of cosmological role seems to be tied to a realist form of objectivism and, in any case, does not seem to be entirely successful in capturing what is at issue. One might seek to avoid these problems by broadening the conception of width of role from

the causal-explanatory to one of width of rational role. The suggestion might be that we can be increasingly confident that judgments in a certain domain are objectively warranted if they ramify through other areas of our rational lives and have a role in rational justifications in these other domains. The greater this ramification, the greater might be our confidence that our judgments are not simply responding to the shadows cast by our own conceptual norms.

This is a very sketchy and rough proposal, but it might seem to have some intuitive appeal. Judgments concerning the properties of physical science clearly ramify very widely, not only throughout a wide variety of fields of our intellectual life, but also practically, by offering us new ways successfully to interact with the physical world around us. Mathematics has a similar width of rational role because it is a central conceptual resource for physical science and a host of other intellectual disciplines. At the other end of the spectrum we can see that more dubiously objective domains fail to have this width of rational role. Judgments of tastiness fail to ramify much beyond the domain of gastronomy. Judgments about the humorous seem to do slightly better, since they do ramify into other areas of our rational lives, for example, our critical judgments about the merits of works of art, literature, music or rhetoric.

This framework might seem to be fairly congenial to the ethical objectivist, since ethical judgments do seem to ramify through large areas of our rational lives. I have already suggested that ethical judgments seem to concern properties with a much wider causal-explanatory role than the ethical anti-objectivist would allow.[17] However, one can

---

[17]Talk of properties here might suggest that I am foreclosing the argument in favour of some form of ethical realism, rather than a wider objectivism. However, in the context of a broadly common sense realism, talk of properties need not suggest that there is anything to which judgments correspond. Thus, in mathematics, one might say that four had the property of being equal to two plus two, but this need not suggest that there is any platonic entity that is being described.

also, with some plausibility, point to the width of the rational role of our ethical concepts. Value concepts feature importantly in historical, economic, sociological, political, aesthetic and anthropological explanations, despite strenuous efforts by reductionists to render these areas value-free. There is also some plausibility in thinking that our very ability to understand our fellow humans at all makes essential appeal to values.[18] Clearly, the way that value concepts ramify throughout areas of our rational lives, other than ethics, is less extensive than the way that mathematical or scientific ones do, but it is, nonetheless, extensive. There will, however, be a serious concern that the ramification in the case of ethical value concepts is more apparent than real.

But there is a fundamental problem with the idea that *width* of rational role is a condition of objectivity. The challenge can be spelt out in this way. Consider the case of magical and astrological explanations. How do these fare under the test that I have suggested? Very badly, one might think, for they fail to give satisfactory causal explanations, which is what they purport to do, and they fail to ramify successfully throughout other areas of our rational lives. This latter claim would be supported by pointing to the widespread revisions that would be necessary to those other, more successful, areas of our rational lives that come into conflict with an understanding of the magical or astrological. It is, for example, hard to see how one could hold on to any remotely sophisticated scientific understanding and still believe in astrological or magical properties. But the point to note is that this response is not one that ultimately relies on *width* of rational role, rather what does the work are the standards of rational criticism

---

[18]See, for example, the essays "Mental Events" and "Psychology as Philosophy", in Davidson (1982a), and Davidson (1982b). Hurley (1989) uses the value-involving nature of interpretation as the basis for an ambitious defence of moral cognitivism.

that determine what is a good explanatory system. For groups who really believe in magic or astrology, such a belief system will have huge ramifications throughout their understanding of themselves and their world. Perhaps a better example is religious belief. Until the Enlightenment, an essentially religious understanding permeated all aspects of people's conception of the world - from science, to morals, to the organisation of social and political structures. Yet despite this widespread ramification, belief in the divine came under sustained and objectivity-threatening challenge. The implication would seem to be that entirely erroneous groups of concepts can ramify extensively and have wide rational roles.[19]

The lesson of this discussion is that width of role, whether the role is seen as a causal one or as a rational, normative one, is not a sufficient condition for objectivity. It also does not seem incoherent to think that a subject matter could be objective, but have very limited width of role of any kind, for example, because it is simply not at all important or significant. This suggests that a width of role requirement is also not necessary for objectivity. What really seem to do all the work in ensuring that a convergence in judgment is objectivity-sustaining are the standards of rational criticism to which these judgments are subjected. This returns us to the question of what form these standards should take if they are to sustain objectivity.

One obvious answer to this question would be to offer a causal-explanatory criterion, but one different from the width of cosmological role requirement. The proposal might be that a judgment that p is objectively warranted if p would feature in the best

---

[19]I do not mean to suggest that some form of religious belief is in error. Rather, there does seem to be a compelling case for maintaining that the sort of all-pervading style of religious belief that existed in pre-Enlightenment minds is untenable.

explanation of what goes on in the world. This would be a causal test, but it need not be seen as dubiously scientistic, because the requisite conception of 'what goes on in the world' could be quite liberal. However, this is best seen as a test for reality, rather than objectivity.

My suspicion is that, if we are interested in the issue of ethical objectivity, we would do best to be satisfied with a test such as David Wiggins's requirement that a vindicatory explanation must be available to show that there is nothing else to think but that p. This requirement is not a causal-explanatory one, but rather asks for an explanation of the normative grounds for the judgment. This test is meant to be a topic-neutral requirement for the objectivity of judgments - it will serve equally well as a test for mathematical, empirical, ethical, or other types of judgment. In each case the kind of vindicatory explanation that is offered will cite different sorts of reasons as providing the normative grounds for the judgment.

The width of cosmological role requirement was initially suggested in order to deal with the intuition that there is something wrong with a standard of objectivity if it allows a cognitive practice just to go on in its own sweet way - if it does not address the possibility that the supposedly independent features to which subjects respond are merely the shadows cast by the practice's own conceptual norms. We can now see that the proposed test does not suffer from this defect because what is required is that "there is nothing else to think for anyone who understands properly what is said by the sentence that says that p".[20] There being nothing else to think is not the same as there being nothing else *for us* to think, rather it means that there is "nothing else to think for anyone at all,

---

[20]Wiggins (1996b), p.272.

anywhere" who understands the relevant concepts.[21] The absolute nature of this requirement means that it will not be compatible with relativization to particular practices. The sorts of consideration that will support a vindicatory explanation will be whatever it takes, from whatever different areas of our rational lives, to establish that there is nothing else to think but that p. Thus, part of showing that there is nothing else to think but that racist behaviour is morally odious might be an appeal to the biological facts about the tiny genetic differences between races.

6. Conclusions

I have tried, in this chapter, to move from a discussion of objectivity at a general level towards an examination of the specific issues surrounding ethical objectivity. Although it is easy to see how the common sense realism, which Wittgenstein seems to advocate, can give vague pointers towards developing a plausible objectivist account of ethics, it is rather less easy to execute this project in detail. We still need to show how ethics might exhibit sufficiently robust standards of correctness for ethical judgments to count as objectively correct in any sense strong enough to be of interest. The robustness of standards of correctness is best seen as lying on a continuum. At one extreme are judgments of the tasty, which are expressed in an assertoric form but make little demand for rational convergence in judgment. At the other end of the spectrum are judgments concerning physical science or mathematics, which demand a high degree of convergence and which also have the reassuring characteristic of ramifying throughout a wide variety of subject matters.

---

[21]Wiggins (1990-1), p.72. See, also, Wiggins (1996a), pp.42-4, which responds to doubts on this score raised by Williams (1996).

This way of thinking about the issue respects the insights we garnered from Wittgenstein, since the conception of objectivity that is involved does not invoke any dubious notion of correspondence to metaphysically privileged objects. I have examined the intuition that objective judgments must be responses to something with wide causal or rational role. In the end, I have concluded that width of role, construed in either way, amounts to a feature that can offer defeasible reassurance that a domain is objective, but it is not a necessary or sufficient condition for objectivity. The reassurance offered is of the following kind. The more judgments in a domain can be seen as responses to things that have a role in explanations, whether causal or normative, in other domains, the less plausible it is to regard the subject matter of these judgments as merely a reflection of our judgmental norms themselves. There seems to be a reasonable prima facie case for saying that ethical judgments can supply some reassurance of this kind. Valuational concepts are central, not just to our ethical lives, but also to many other rational domains that treat humans as practically rational, thinking beings - history, economics, sociology, anthropology, political science, and so on. But this does not tell us anything about the status of particular ethical judgments, and this reassurance does not amount to a sufficient condition for objectivity. Ultimately, we returned to the idea that the objective correctness of a judgment that p is a question of whether one comes to believe p because reason demands that there is nothing else to think but that p. The intuition that a cognitive practice cannot secure its objectivity simply by insulating itself from other practices is respected. However, this is not through a width of role requirement, but through the requirement that there be nothing else to think for anyone thinking properly about the subject. The question is whether ethics displays norms of reason that are sufficiently robust to issue in judgments with this kind of status.

So the task that remains is to develop an account of ethical objectivity that gives us a substantial standard of rational criticism. The account will also need plausibly to explain how we come to reach ethical judgments that are consciously aimed at meeting these standards. To pursue this task, I shall first examine in some detail in the next chapter John McDowell's attempt to argue that Wittgenstein's rule-following considerations entail a modest form of neo-Aristotelian virtue-centred ethics. The problem with this approach is that the very means by which it seeks to secure ethical objectivity undermine any confidence that we might have that rational criticism in ethics can support absolute, objectivity-sustaining vindicatory explanations of any convergence in judgment. It also fails, I shall argue, to give a plausible account of ethical deliberation. I then go on, in Chapter 6, to argue for a form of cognitivism that uses as its mechanism for rational criticism Kant's notion of the Categorical Imperative. This, I shall argue, overcomes the sort of weakness with respect to objectivity that I find in McDowell's approach. It also offers a rather more plausible model of ethical deliberation than does McDowell's virtue-centred theory.

Chapter 5

AGAINST ONE VIRTUE-CENTRED THEORY OF ETHICAL DELIBERATION

1. One Model of Ethical Realism

The challenge that we were left with at the end of the previous chapter was to give some account of the standards of rational criticism that are available in ethics and to show that these are substantial enough to yield objectivity for ethical judgments. There would be something seriously lacking from a meta-ethical theory that simply provided a theory of value that treats the world of human actions as a set of scenes to be described and anatomized. Rather, such a theory must also show how valuations can guide and foster these actions. It follows from this that a meta-ethical theory that is objectivist cannot simply say that certain ethical descriptions are such as to apply truly or falsely, but must also show that, in some way, those human actions can be guided by reasons that are objective. So, in addition to an account of the objectivity of ethical judgment, there must be an account of practical reason that says something about how these judgments guide our actions.

I identified the general framework within which we should operate as a non-reductionist form of response-dependent objectivism. In this chapter, I shall examine a particularly restrictive version of this approach, to be found in the ethical writings of John McDowell. In a series of influential articles over the last twenty years, McDowell has sought to advance a form of ethical objectivism that combines a rejection of a scientistic - or 'absolute'- conception of reality with an appeal to a characteristically Aristotelian

conception of ethical judgment.[1] Whilst deeply sympathetic to the first element of this project, I find the second far less convincing, making, as I believe it does, objectivism about ethical judgment less plausible than it in fact is.

I focus on McDowell's work for two reasons. The first is that the ethical theory that he has developed is one of the most influential non-reductionist forms of response-dependent ethical objectivism, or realism, as he prefers to call it. Of those who seek to develop response-dependent ethical objectivism by advancing a virtue-centred ethics, his theory offers one of the most clearly worked out modern accounts of the role of virtues in practical deliberation. The second, and more important, reason is that he explicitly follows a very similar route to the one that we have pursued, for he sees Wittgenstein's discussion of rule-following as offering some important lessons for philosophical ethics. One of the key lessons that he takes from Wittgenstein is that the rule-following considerations put substantive constraints on the way in which practical rationality should be conceptualized. If these restrictions are in good order, then there is no need to develop and, indeed, no possibility of developing the sort of ethical objectivism that I wish to. Thus, it is important that I engage with McDowell's arguments to show that his highly restrictive approach is not compulsory simply because one accepts Wittgenstein's arguments about objectivity, and that there is also a need to go beyond his theory of practical deliberation. In order to reach the latter conclusion, I shall argue first that his theory is so restrictive that it lacks the conceptual resources to offer a genuine form of ethical objectivity, and secondly that the model of practical deliberation that he offers is seriously flawed.

---

[1]See, in particular, McDowell (1978), (1979), (1980), (1981a), (1995a) and (1995c).

McDowell has advocated the view that: (a) objectivity is not a matter of judgments within a given domain matching up to external, reductionistically-defined standards; (b) judgments in the domain of ethics can satisfy a properly construed standard of objectivity; (c) what determines correct ethical judgment is the judgment of competent, adult virtuous agents, when subjected to appropriate internal criticism; and, (d) in determining these standards of internal criticism, the particular sensitivities of virtuous agents are conceptually prior to any universal principles - in other words, these standards are uncodifiable. In the first half of the chapter, my concern is with the effect that this last claim - what I shall dub the 'Uncodifiability Thesis' - has on the rest of the enterprise of developing a response-dependent form of ethical objectivism. I shall argue that McDowell's grounds for finding some intimate connection between the Wittgensteinian reasons for advocating a non-reductionist response-dependent ethical objectivism and those for adopting the Uncodifiability Thesis are mistaken. Worse still, I shall argue that, not only is the Uncodifiability Thesis not entailed by the adoption of a non-reductionist, response-dependent ethical objectivism, it actually fatally disables it. This is because it limits the conceptual resources available to the ethical objectivist to such an extent that it is not possible to maintain the distinction between subjects judging incorrectly and their simply making judgments involving different concepts. Without this distinction, there can be no real sense in which subjects judge correctly or incorrectly, and hence no objectivity.

So, the first half of the chapter deals with disentangling and analysing the metaphysical aspects of McDowell's project. A response-dependent objectivism needs to have at its heart a theory of practical rationality. I will have argued that McDowell's favoured neo-Aristotelian, virtue-centred theory is unsupported by his arguments for it and fails to deliver genuine objectivity for ethical judgment. In the second half of the

chapter, I examine McDowell's account of practical deliberation in order to argue that, ultimately, this offers no plausible model of such deliberation. This is a serious lacuna for a response-dependent conception of ethical objectivism. Once again it will be seen to be the debilitating effect of the Uncodifiability Thesis that produces this result. McDowell claims that his views on practical deliberation are a development of Aristotle's. I should note that it is not within the scope of my project to form a judgment on whether or not they really are a correct interpretation of Aristotle. Maybe an Aristotle different from McDowell's provides all the answers - I do not know. But I do suspect that an attempt to develop a virtue-centred ethics - such as McDowell's - that avoids grounding the virtues in an extra-moral conception of human excellence is the one that would stand the best chance of saying something serious and useful to us today.

2. Uncodifiability and the Limits of Internal Criticism

I begin by giving a very brief and schematic presentation of the theory of response-dependent ethical realism that McDowell develops. Two theses are essential to McDowell's position but, I shall argue, they are ultimately in tension in a way that is fatal to his objectivist aims. Fortunately for objectivists, one of the theses will be found to be uncompelling.

The first thesis is:

The Internal Criticism Thesis. The objectivity of a particular domain of enquiry is not a matter of its having an external grounding independent of the reactions of

> competent subjects, but of judgments within it withstanding criticism 'from inside'.[2]

Throughout the course of his work on ethics, McDowell has defended this idea and has done so in at least four closely related ways: he has argued that Wittgenstein's consideration of rule-following supports a rejection of the idea that normativity is underpinned by some non-normative structure, available for scrutiny from a perspective external to all normative practices;[3] he has argued against the attractiveness of equating objectivity with the possession of primary qualities, or with accessibility to a 'view from nowhere';[4] he has developed an exegesis of Aristotle that sees him as rejecting the need for a response-independent conception of human flourishing or eudaimonia;[5] and, recently, he has tended to rest his case on an appeal to Kant's views on spontaneity, shorn of their idealism.[6] But the point has always been to maintain something like the Internal Criticism Thesis.

I shall focus on the Wittgensteinian route. In the first place, this is because my interest is in seeing what lessons we can draw from Wittgenstein's discussion of

---

[2]Although McDowell never states this thesis in so many words, it is clear that he subscribes to it. See McDowell (1980), (1995a) and (1995c) for his account of how internal criticism is to work. For another way of bringing out the key claims that McDowell makes, see Charles (1995), pp.138-41. The Internal Criticism Thesis roughly corresponds to a combination of Charles's claims [B] "Objectivity Secured by the Withstanding of Internal Criticism" (p.138) and [C] "No Reaction-Independent Theory of the Human Good" (pp.138-9).

[3]See McDowell (1979) and (1981a) for this approach.

[4]See McDowell (1983), (1985b) and (1986).

[5]See McDowell (1980) and (1995a). This is, perhaps, less an independent defence of the thesis and more an attempt to attribute it to Aristotle and give some suggestions as to how it might work.

[6]See McDowell (1995c), and also (1994).

objectivity for ethical objectivity in particular. Secondly, it is here that the connection with the Uncodifiability Thesis is most explicitly argued for, and it is this thesis that has the most dramatic effect on how we are to understand practical rationality.

It is the primary aim of McDowell's paper "Non-Cognitivism and Rule-Following"[7] to use Wittgenstein's ideas about rule-following as a weapon with which to attack the basic metaphysical assumptions that he thinks are held by ethical non-cognitivists. These metaphysical assumptions, which exclude values from reality, are that reality is restricted to things 'as they are in themselves'. In other words, the only things that have genuine reality are, roughly, primary qualities, describable independently of any particular perspective. The non-cognitivist claim is that values cannot count as possessing this kind of reality, because our value judgments are not responses to things conceived of in this way. Rather, we make judgments about values simply because our nature or nurture has given us certain affective propensities that lead us to make certain evaluative classifications.

Non-cognitivists are confident that they can give an adequate account of ethical judgment that makes no appeal to values in the world, and McDowell summarises this account thus:

> Typically, non-cognitivists hold that when we feel impelled to ascribe value to something, what is actually happening can be disentangled into two components. Competence with an evaluative concept involves, first, a sensitivity to a certain aspect of the world as it really is (as it is independently of value experience), and, second, a propensity to a certain attitude - a non-cognitive state which constitutes the special perspective from which items in the world seem endowed with the value in question.[8]

---

[7]McDowell (1981a).

[8]McDowell (1981a), p.143.

On the non-cognitivists' account of ethical judgment, an evaluation has two elements. First, there is a purely cognitive belief about the world, as described in respectable primary quality, perspective-free terms. Secondly, there is an affective element, such as a desire or pro-attitude, that is projected onto the feature picked out by the belief. The crucial move that non-cognitivists make is to claim that this affective element is like any other appetitive state because it can be characterized independently of whatever features of the world that it is projected onto. McDowell puts the implications for reason explanation of this claim thus:

> The rationality that a reason explanation reveals in an action it explains ought, if the explanation is a good one, to be genuinely there: that is, recognisable from an objective standpoint, conceived in terms of the notion of the view from sideways on ...[A]nd if we think someone's possession of the desires in question can be recognized from a standpoint external to the agent's moral outlook, then it might seem that those desires would confer an obvious rationality, recognizable from that objective standpoint, on actions undertaken with a view to gratifying them.[9]

However, the cogency of this disentangling manoeuvre, which the non-cognitivist suggests, is not beyond dispute:

> If the disentangling manoeuvre is always possible, that implies that the extension of the associated term, as it would be used by someone who belongs to the [moral] community, could be mastered independently of the special concerns which, in the community, would show themselves in admiration or emulation of actions seen as falling under the concept... But is it at all plausible that this singling out can always be brought off?[10]

This might seem to be a denial that the evaluative supervenes on the non-evaluative. However, it is important to note that this is not what McDowell's denial of the disentangling manoeuvre entails, because non-cognitivists are making a more ambitious claim. For non-cognitivists, what the evaluative classification supervenes upon must form

---

[9]McDowell (1981a), p.155.

[10]McDowell (1981a), p.144.

a genuine kind, describable in terms of the level supervened upon. However, it seems permissible to entertain doubts about this claim:

> [H]owever long a list we give of items to which a supervening term applies, described in terms of the level supervened upon, there may be no way, expressible at the level supervened upon, of grouping just such items together...Understanding why just those things belong together may essentially require understanding the supervening term.[11]

The non-cognitivist's assumption is that it will always be possible, in principle, to identify such kinds at the base level. But McDowell suggests that this assumption is based on the tempting picture of concept application that is offered by metaphysical or philosophical realism. Thus, the application of a concept is objectively warranted by a subject's latching on to some platonic 'rule-as-rail' or acting out of some ideal, reductively defined disposition. Under either interpretation, the aspiration is that the grounds for warranted use of a concept can be characterised in a manner that is non-normative and independent of the responses of the participants in a particular concept-deploying practice. McDowell characterizes the non-cognitivist's hope thus:

> [I]t is in principle discernable, from a standpoint independent of the responses that characterize a participant, that a series of correct moves in the practice is really a case of going on doing the same thing.[12]

For reasons similar to those discussed in Chapter 3.2-3.5, McDowell sees Wittgenstein's examination of rule-following as demonstrating that it is incoherent to demand that the normativity of our rational practices be given a grounding that is from a 'side-on' perspective external to these practices. As we saw in Chapter 3, in Wittgenstein's opinion we should recognise that the normativity of meaning depends on the existence of a background of shared custom - shared sensitivities, routes of interest

---

[11]McDowell (1981a), p.145.

[12]McDowell (1981a), p.146.

and basic behaviour. Instead of grounding the objectivity of our use of concepts more securely, the attempt to demand or to offer justifications from outside the 'whirl of organism' of our rational practice simply results in the loss of normative constraints.[13] A radically independent perspective cannot be what grounds the objective deployment of concepts, because such a perspective is simply not there to be occupied. If an appeal to this totally external perspective is incoherent, then, McDowell claims, it is no slight on the objectivity of ethical values if they are not to be found from that perspective.

This puts into play the first of McDowell's theses, along with a sketch of his reasons for holding it. For reasons with which we should be familiar from the previous two chapters, McDowell has argued that what sustains objectivity is criticism and rational reflection from within our conceptual practices. But, as we have seen, there is plenty of scope to ask what form this criticism should take. McDowell's answer to this question is the second key element of his theory that I want to discuss. This is (what I shall call):

The Uncodifiability Thesis. The individual reactions of competent adult ethical agents are conceptually and explanatorily primary in determining how the process of internal criticism is to function. Principles or rules can, at best, have a purely secondary role.[14]

We can approach a discussion of this thesis by considering a debate about the basic purpose of ethics, about how we tackle the question of how one should act. As McDowell

---

[13]See the discussion of meaning irrealism in Chapter 3.5 above.

[14]Again this is a distillation of what McDowell says, rather than something he explicitly states. This thesis corresponds roughly to Charles's [E] "The Centrality of Case Specific Reactions: Difficulties in the Very Notion of a General Theory of the Human Good" (Charles (1995), p.140).

puts it, on one approach to this question, "the primary topic of ethics is the concept of right conduct, and the nature and justification of principles of behaviour".[15] A consequence of this view is that the concept of virtue or moral character will be of secondary interest, for virtue will simply be definable (very crudely) as a state of character that issues in (independently-defined) right conduct. In contrast to this approach, McDowell wants to support a very different answer to the question, "How should one live?". On this view the question is:

> [N]ecessarily approached via the notion of a virtuous person. A conception of right conduct is grasped, as it were, from the inside out.[16]

The argument for this claim is to the fore in his other paper that covers the implications of Wittgenstein's rule-following considerations for ethical objectivity, "Virtue and Reason".[17] In addition to deploying these ideas to undermine non-cognitivism, McDowell also wants to use them to defend the implications of the Uncodifiability Thesis for the way in which ethics is to be conducted. His aim is to challenge a conception of rationality that is inhospitable to the Uncodifiability Thesis. But there also seems to be another, more ambitious claim that he sees his appeal to the rule-following considerations as supporting. This is the idea that ethical theories that give a conceptually primary role to principles of right conduct are in breach of the "thesis of uncodifiability" that he finds in the rule-following considerations, and hence only a virtue-based theory of ethics will be legitimate.

---

[15]McDowell (1979), p.331.

[16]McDowell (1979), p.331.

[17]McDowell (1979).

We can find evidence for this more ambitious claim in passages where McDowell attacks the idea that we can give a ranking of the concerns that feature in a virtuous person's conception of the good life:

> But uncodifiability rules out laying down such general rankings in advance of all the predicaments with which life may confront one.[18]

This suggests that any codification of correct conduct will be strictly limited, a conclusion that is explicitly brought out towards the end of the paper:

> If the question "How should one live?" could be given a direct answer in universal terms, the concept of virtue would have only a secondary place in moral philosophy. But the thesis of uncodifiability excludes a head-on approach to the question whose urgency gives ethics its interest. Occasion by occasion, one knows what to do, if one does, not by applying universal principles but by being a certain sort of person: one who sees situations in a certain distinctive way. And there is no dislodging, from the central position they occupy in the ethical reflection of Plato and Aristotle, questions about the nature and ...the acquisition of virtue.[19]

It is not beyond question that this kind of substantive claim about how we are to understand ethics can be derived from the rule-following considerations, which are, on my interpretation, an attack on a philosophical misunderstanding of the nature of objectivity. McDowell's more ambitious claim would be in order only if ethical theories that appeal to rules or principles of right conduct are guilty of the sort of codification that Wittgenstein attacks. But they are not - or so I shall argue in the next section.

Two key features of McDowell's moral realism are encapsulated by the two theses that I have outlined. McDowell seems to see the first as somehow entailing the second. The idea of principles as fundamental in ethical deliberation supposedly involves an appeal to a conception of rules as external to and grounding normative practices. He takes the Internal Criticism Thesis, and the Wittgensteinian case for it, to have undermined this

---

[18]McDowell (1979), p.344.

[19]McDowell (1979), p.347.

conception. The result of conjoining these two theses is a form of moral realism that is tied to a highly particular, theoretically very modest - even restrictive - version of Aristotelianism concerning ethical deliberation.

3. Codification and Deliberative Rules

So, McDowell seems to think that, once we reject the idea that objectivity is a matter of validation by factors outside ethical practice, we must also reject the conceptual primacy of principles in ethical deliberation. However, there seems to be an obvious response to this Uncodifiability Thesis. In Chapter 3.6 and 3.7, it was argued that it is important not to regard Wittgenstein as offering a reductionist communitarian account of objectivity, because such an account would inevitably seem unsatisfactory. Instead, I argued that Wittgenstein is best seen as some kind of common sense realist who does not have to offer a constructive account of meaning, precisely because he regards the demand for such an account as hopelessly confused. The rule-following considerations are aimed at attacking confused metaphysical pictures of what it is to understand a concept correctly, namely metaphysical or platonic realism on the one hand and meaning scepticism on the other. On this interpretation, it is quite in order to use the rule-following considerations to attack a confused metaphysical picture that underlies a meta-ethical theory about ethical objectivity. Thus, we can use the rule-following considerations, as McDowell does, to attack the metaphysical assumptions that underpin ethical non-cognitivism. We can therefore reject as incoherent the demand that objectivity involve, as one might put it, codification 'all the way down'. The question is one of how much more we can glean from the rule-following considerations about the nature of ethics or practical rationality.

One question that one might want to ask is whether ethical theories that do give a conceptually primary role to principles or rules always involve some form of illegitimate codification of a metaphysically ambitious kind. In a short while, I shall briefly introduce two views, according to which principles might play a central role but which seem to involve no such metaphysical over-ambition. Before considering these, I want to say a few more things at a fairly abstract level.

One of the key points of the common sense realism that I attributed to Wittgenstein was an aspiration to rehabilitate our ordinary, non-philosophical ways of justifying the use of concepts. The rule-following considerations are simply meant to 'factor through' and leave everything, substantially speaking, unchanged. The only thing that is changed is our tendency to demand some kind of philosophical realist justification for our use of concepts. In effect, the rule-following considerations support a very much weaker version of the Uncodifiability Thesis than the one that McDowell relies upon, namely something like:

> *The Weak Uncodifiability Thesis.* The objectively correct application of a concept cannot be codified in terms that are entirely external to norm-governed conceptual practices. The objectively correct application of a concept is the application of the concept in a way that is justified by good reasons (whatever these may be).

This thesis seems to be an accurate reflection of a common sense realism about objectivity. It does not foreclose the issue of what might, or might not, count as good reasons within a practice - it might be a good reason to point to sensitivities and perceptions, but it might not; it might be a good reason to point to a rule or principle, but it might not.

On this view, we can carry on using mathematical, semantic, or logical rules or norms in a perfectly proper way. We must just drop the idea that this usage is any less in order if the rules or norms do not reflect features of a platonic realm (as they do not). Similarly, we might think that the 'deductive paradigm' could be in perfectly good order as a model of practical rationality. In all these cases, the existence of shared sensitivities, basic reactions, customs and practices will be essential for the objectively correct use of concepts, but it is not legitimate actually to appeal to these in giving a justification for one's conduct. In order to give a justification, one must offer a reason, and, for anything that has yet been said, this might involve essential reference to a rule or principle.

Perhaps the reason that McDowell feels that he can appeal to a strong form of the Uncodifiability Thesis is that he is tempted by a communitarian conception of meaning. This is not a position that he specifically endorses in either of the two papers that we have been discussing, and in later papers he explicitly rejects it as an appropriate understanding of Wittgenstein's discussion of rule-following.[20] I have argued that this is an interpretation of Wittgenstein that should be rejected, both because of its intrinsic weakness and because of its exegetical inaccuracy. If McDowell does not rest his strong Uncodifiability Thesis on such an interpretation, then it is hard to see why he feels justified in endorsing it, rather than something like the Weak Uncodifiability Thesis.

I have suggested that even the decision-procedure model of the role of principles in practical rationality might seem untouched by McDowell's argument, when this is seen as establishing the Weak Uncodifiablity Thesis, rather than the stronger version. In spite of anything said so far, a codification of rationality, practical or otherwise, might be in

---

[20]See McDowell (1984) and (1992).

order, provided it takes place within our rational practices. The opponent of the deductive paradigm will need to advance some other kind of argument if the deductive paradigm, understood in a way that does not flout the Weak Uncodifiability Thesis, is to be undermined as a model of rational decision-making. There is a tactical difficulty that faces anyone attempting this task, namely that the supposed uncodifiability of rationality will not, by its very nature, be susceptible to conclusive proof. Two ways for champions of uncodifiability to move forward are suggested by William Child.[21] In the first place, they might challenge the reductionist to give reasons for holding that rationality is codifiable, and then knock these down. This approach could be supplemented by unearthing, as McDowell attempts to, the reasons why one might be tempted to adopt such a position in the first place. The second strategy would be to offer positive considerations, which do not amount to a proof, to illustrate why there is no possibility of codification. Examples of such considerations might be the claim that no plausible codification has ever been proposed, or an appeal to the fact that codifiability is fatally undermined by the lack of codifiable principles to adjudicate between conflicting theories or criteria.

It seems plausible to suppose that there is a very heavy burden of proof resting upon the shoulders of those who advocate the codifiability of rationality in a strong way that would flout the Weak Uncodifiability Thesis. However, the points that Child raises do not seem to touch the claim that in certain areas of rationality there is codification and that this is of the weak sort compatible with our Weak Uncodifiability Thesis, and this is accepting of the essential role of subjects' sensitivities and judgment. Indeed, there seems to be a heavy onus of proof on those who deny that certain, fairly large areas of rationality

---

[21]Child (1993). His primary concern is with the uncodifiability of rationality in general and its implications for reductionism about the mental.

are best understood in a way that makes essential use of codification of an innocuous sort. Yet it is this position that is denied by someone who appeals to the uncodifiability of rationality as part of a positive defence of the central role of virtue in our understanding of ethics.

Let us now turn to a different line of attack on the strong Uncodifiability Thesis. We can accept that metaphysically ambitious codification of the type hankered after by platonists about meaning or reductionists about the mental is incoherent. We can even accept that the 'deductive paradigm' is flawed. However, there is still an important question about whether there are not other ways in which rules or principles could play a central role in our understanding of an area of rationality but not fall foul of this attack. If there are, then the case for a virtue-centred ethics has not been advanced. It would appear that, in McDowell's view, rules would have to be part of a quasi-mechanistic decision-procedure if they are to have a central role in practical rationality, and that for them to play this role, they have to be conceived of as codifiable in a strong sense. The picture that this model suggests is one of a subject slotting into the mechanism the particular features of the situation and waiting for an answer to pop out. As I have accepted, McDowell seems justified in opposing this model if the codification is to go 'all the way down'. I want briefly to introduce two alternative ways of understanding rules that seem to remain untouched by McDowell's arguments. These are (a) the concept of rules as constitutive of practices, rather than summaries of correct behaviour,[22] and (b) principles of rationality as negative norms for the recursive criticism of judgments. I introduce these simply to illustrate other ways in which, prima facie, rules and principles

---

[22]For the classic exposition of the distinction between summary and constitutive rules, see Rawls (1955).

might seem to escape the arguments for the Uncodifiability Thesis. The first of these proposals I shall not pursue any further, the second I shall pick up and expand upon in Chapter 6.

(a) *Constitutive Rules*

I shall introduce the concept of constitutive rules by considering the example of the rules that might appear in a cricket manual. In such a manual, we might find a rule such as, "Avoid playing across the line of the ball". This is a rule that acts as a summary or generalization of the behaviour that generally produces the best results - it is a rule of thumb. Most of the time a batsman will have the best chance of hitting the ball in a controlled way if he or she strikes in line with the flight of the ball. However, players with a particularly good eye for the ball might find that they can ignore the rule and score more runs by hitting across the line of the ball. As Aristotle says, and McDowell agrees, such a rule only holds for the most part.[23] The practice of playing cricket is conceptually prior to the formulation of this rule. Even if no one had codified this piece of advice, the playing of cricket would have carried on essentially unchanged (but for rather more wickets falling).[24]

In contrast to these summary rules are rules that are constitutive of a practice. Examples of such rules are rules that specify that a player is deemed out if they are caught, bowled or found to be leg before wicket. With this sort of rule, the practice of playing cricket is partially defined by the existence of the rule, and hence the rule is logically prior

---

[23]See Aristotle (1985), I.3.

[24]Indeed, followers of the English cricket team might be forgiven for thinking that this is how cricket *is* played.

to action in accordance with it. Although there could be changes to the details of such rules (as there have often been with the leg before wicket rule), significant changes to centrally important rules would mean that a different game was being played. Because of their constitutive role and their generality, such norms have no pretension to act as codifications of correct behaviour that might figure in a decision-producing mechanism. Because of their constitutive nature, they cannot coherently be regarded as holding only for the most part. We can even agree with Child when he writes:

> To say that rationality is uncodifiable is not necessarily to say that there can be absolutely no true, exceptionless principles of rationality. But if there are any such principles, they will not be such as to deliver a detailed answer to every question (or, more ambitiously, to any question) of the form, What should I do, or believe, in these circumstances?[25]

Constitutive norms do not have the function of giving answers, nor, *a fortiori*, detailed answers, in some decision-theoretic manner. Their role is to constitute what an appropriate answer would look like. This would seem to suggest that constitutive rules or norms might be compatible with the Weak Uncodifiability Thesis.

Our cricket example can also shed light on the relationship between a constitutive rule and the sensitivities of the agents applying the rule. Clearly, no one thinks that the rules of cricket are self-deploying in some mysterious way. The application of such rules requires considerable expertise. The judgment that a batsman has been caught will require the ability very accurately to perceive the flight of a fast moving ball and its interaction with the bat, the fielder's hands and other obstacles. The decision on whether a batsman is out leg before wicket requires, not only the ability to analyse the flight of the ball, but also to project what the flight of the ball would have been had it not hit the batsman. But the essential role of such sensitivities and cognitive abilities does not undermine the fact

---

[25]Child (1993), p.219. I might, however, want to take issue with Child's parenthetic remark.

that the rule is logically and conceptually primary. If there is a dispute about an umpire's decision, the dispute will be about whether the umpire has correctly judged according to the rule. The dispute will not be fully comprehensible except as a dispute about the application of the rule. Any appeal to the sensitivities or cognitive skills of the umpire will be of merely derivative significance.

Clearly, our cricket example is a very simplistic one. But it does highlight the way that we can think of a rule as being conceptually primary (so that we can only come to an understanding of the subject by reference to the rule), without this importing any metaphysically weird notion of codification all the way down. Can we apply this model of rules to something a bit more complex than a game of cricket? One rather more philosophical example of such norms are the norms of rationality that Donald Davidson believes are constitutive of the mental. These norms have such a role precisely because, claims Davidson, they are uncodifiable and "have no echo in physical theory".[26] The centrality of these norms is envisaged as being precisely a consequence of their uncodifiability. How might this idea extend to ethics? One thought might be that constitutive rules play a role in our understanding of particular moral institutions, such as punishment or the making and keeping of promises. This was the context in which John Rawls originally highlighted the distinction between constitutive and summary rules. More ambitiously, someone might try to argue that the satisfaction of a Kantian Categorical Imperative test, the maximization of utility or acting in accordance with a certain set of duties is constitutive of morally right conduct.

---

[26]Davidson (1982a), p.231. For more on this idea see the essays "Mental Events", "The Material Mind", and "Psychology as Philosophy", in Davidson (1982a). McDowell (1985a) provides a more detailed development of the normative role of these constitutive norms and the reasons for thinking that they cannot be reflected in the physical world.

There are, of course, clear disanalogies between the cases of constitutive norms of cricket and any putative ethical ones, beyond merely the extreme simplicity of cricket. Two obvious ones are that the cricket rules are explicitly codified and that they are optional.[27] Analogies will always carry with them disanalogies, so the question is whether the disanalogies are significant.

The first disanalogy seems pretty irrelevant to the point being made. I introduced constitutive rules or norms as an example of a kind of rule that was not decision-theoretic, had no metaphysical ambitions, and was compatible with a recognition of the importance of subjects' sensitivities and cognitive skills, but was, nonetheless, explanatorily and conceptually primary. Rules of a game exemplify this, but there is no reason to think that these features necessarily require the rule to be explicitly codified. Every possible application of the leg before wicket rule is not written down - it requires the judgment and understanding of the umpire to implement the rule. The example of constitutive norms of rationality illustrates this point even more forcibly.

Moving on to the second objection, it is indeed true that it is optional whether we participate in the practice of playing cricket, but if we do participate it is not optional whether or not we obey the laws of cricket. The issue of optionality seems to be orthogonal to that of the relative merits of virtues and principles in understanding ethical deliberation. If one accepts a certain practice, it is not optional whether one obeys its constitutive norms, but it might be optional whether one participates in the practice in the first place. Similarly, if one has a certain disposition of character, then certain types of action will not seem live options, but it might be optional whether or not one acquires that

---

[27] I am grateful to Bill Pollard for raising this objection.

character trait. Optionality seems to be a feature of the subject matter, not of norms themselves. Games seem highly optional; morality or, *a fortiori*, rationality are less easily seen as optional, because, arguably, it becomes unclear what it is to opt out. Thus, the optional nature of the norms of cricket seems to be a feature of cricket, rather than of norms themselves.

Clearly, an awful lot of work would need to be done to justify the claim that such norms are constitutive of morally right conduct. Ultimately, I do not want to pursue this kind of proposal because there will be a tendency for it to appear to involve a form of decision-theoretic rule, rather than a genuinely constitutive norm. However, my aim here is not to propose a fully-fledged conception of ethical rules, but merely to point out that there seem to be conceptions of rules, and their role in ethics and rationality more generally, that are not obviously targeted by McDowell's arguments. The debate about whether there are constitutive rules or principles in ethics will focus on the question of whether our understanding of certain types of actions and their choiceworthiness makes essential reference to (metaphysically unambitious) rules or principles. For example, when we ask why we have reason to obey a promise, will the explanation give an unavoidable conceptual role to the rule-involving notion of a promise, or will it essentially end up making reference to what a virtuous person would do? If we accept the Weak Uncodifiability Thesis and reject as unsupported the original Uncodifiability Thesis, then this argument is going to have to be battled out case by case. The advocate of a virtue-centred ethics cannot use the Uncodifiability Thesis to argue for the primacy of virtue-centred explanations without begging the question. Although I do not want to pursue the

route offered by constitutive rules, I suspect that the case for the explanatory utility of the appeal to virtues may well have been overstated.[28]

So, I have suggested that a conception of rules as constitutive offers an understanding of norms that are conceptually primary and provide the only route to an understanding of the actions that are required by them. But these rules do not make any claim to be intelligible from a side-on or external perspective. Such norms seem to be features internal to our conceptual practices, and so might be seen as offering an alternative way of approaching the notion of right conduct "from the inside out", as McDowell puts it, without importing a virtue-centred conception.[29]

(b) *Recursive Norms of Rational Criticism*

As a second illustration of how one might think of rules or norms as not involving a dubious codification, I want to introduce a brief advertisement for what is to come in the next chapter. Related to constitutive norms are negative recursive principles that are used in the criticism of judgments. Because of their negative and purely recursive nature these also seem to escape the threat of the Weak Uncodifiability Thesis. Let me illustrate this proposal with a very brief and rather schematic example (the idea will be much more fully explored in the next chapter).

Onora O'Neill has argued that Kant envisaged the critique of reason as a practical task governed by the negative regulative maxim that "any principles of thinking and acting that can have authority cannot enjoin principles on which some members of a plurality

---

[28]I have more to say on in this Section 6 below.

[29]See McDowell (1979), p.331, for the quoted phrase.

cannot (not 'would not'!) act".[30] She claims that this principle effectively comes to the Categorical Imperative (CI). O'Neill suggests, by a process of elimination, that we have to see the authority of reason as self-imposed and its vindication as a recursive process relying on no predetermined procedures or positive algorithms that would yield supposedly completely determinate results. Algorithms are empty on their own and require premises to be fed into them if they are to yield results. These premises will either be groundless and unvindicated, or else be open to vindication by reason, and thus reason cannot be conceived of as entirely algorithmic. As O'Neill puts it, "Thought and action can at most be constrained, not fully determined by principles of reason".[31] Indeed, the principles of reason must, themselves, be provisional and open to criticism by reason, or else they too would be unvindicated, thus stymying the project of vindicating reason. So, principles of reason have to be purely negative norms, rather than positive algorithms, that can be recursively applied in order to vindicate themselves and that leave matters largely indeterminate.

Principles such as these seem to escape any threat posed by the Weak Uncodifiability Thesis. The CI test, on any of its formulations, would be a strange sort of algorithm, if algorithms are to be regarded as yielding fully determinate results. It is highly abstract and purely negative in nature, and (as the hackneyed charge of vacuous formalism attempts to show), when taken on its own, its action-guiding content is fairly non-existent. But, as its extreme abstractness suggests, Kant had no intention of using it as an algorithm. Indeed, such a move would have been utterly alien to his project because he

---

[30]O'Neill (1989), p.20.

[31]O'Neill (1989), p.19.

recognized that, since algorithms are not self-deploying, they cannot play a fundamental role in reason:

> General logic contains, and can contain, no rules for judgment... If it sought to give general instructions how we are to subsume under these rules, that is, to distinguish whether something does or does not come under them, that could only be by means of another rule. This in turn for the very reason that it is a rule, again demands guidance from judgment. And thus it appears that, though understanding is capable of being instructed, and of being equipped with rules, judgment is a peculiar talent which can be practised only, and cannot be taught.[32]

This seems to be as explicit a rejection of the idea that rationality should be fully codifiable as one could want, and yet principles of reason are central to this non-foundationalist account of reason. Clearly, a lot more needs to be said to flesh out this proposal and this is something that I shall do in the next chapter, but I hope that this gives a taste of another conception of norms that seems to involve no illegitimate codification.

So, I conclude that the positive case for the Uncodifiability Thesis, as opposed to the Weak Uncodifiability Thesis, has not been sustained. The Weak Codifiability Thesis does seem to be supported by Wittgenstein's rule-following considerations, but it is not clear that it even rules out a decision-theoretic conception of rules. I have also introduced two alternative conceptions of rules that certainly seem to have no problem with this thesis. The original Uncodifiability Thesis goes much further than is justified by the rule-following considerations and stands in need of justification. The problem with rules or principles of rationality arises only when these acquire a particular sort of exaggerated metaphysical pretension. But an argument that operates in the abstract realm of the metaphysics of meaning cannot be used to argue for anything but the most general background theoretical

---

[32]Kant (1929), A133/B172.

constraints in the mundane world of metaphysically unambitious ethical rules and principles.

4. Normativity and Substantive Disagreement

So far we have been engaged with the question of whether McDowell's Uncodifiability Thesis is really justified and, if it has not been justified, whether there are other ways of understanding ethical norms or principles that do not fall prey to any Wittgensteinian considerations. It is compatible with what has been argued so far that the virtue-centred approach advocated by McDowell offers a coherent and successful account of how ethical judgments might be objective. My aim in this section is to challenge the idea that McDowell's account of ethical practice can achieve this goal.

An obvious reaction of someone who is unfriendly towards non-reductionist, response-dependent theories of objectivity, when faced with the combination of McDowell's two theses, might be an objection like this:

> *The Which Ethics? Objection.* The point of ethical objectivism is that it amounts to some kind of achievement to come to the right answer. But McDowell's insistence on the sufficiency of criticism within a practice means that this hope is empty because other, non-virtuous ways of judging might be equally capable of sustaining this type of criticism. So what is so special about the virtuous?

The thought is that some other group of subjects could respond to situations in a way that differs from that in which the virtuous respond and that this group could constitute a coherent practice of their own. This seems to conflict with the demand that, if there is to

be ethical objectivity, then there cannot be a constant danger that apparently conflicting answers count as ethically correct. Yet the Internal Criticism Thesis, when combined with the Uncodifiability Thesis, might seem to bring just this danger. The superiority of the virtuous way of judging cannot (because of the Internal Criticism Thesis) be demonstrated by appeal to external goods or external validation. The Uncodifiability Thesis means that this condition is further tightened because the sensitivities and character of the virtuous will be essentially invoked in any explanation of this superiority.

Before canvassing, and then rejecting, a response to this objection I want to add a second objection that is a corollary of it:

The Substantive Disagreement Objection. Objectivity requires the possibility of drawing a distinction between judging wrongly and judging differently (that is, using different concepts). It is doubtful whether this kind of distinction can operate if McDowell's combination of theses applies, because persistently unvirtuous judgers would count as doing something other than ethics, rather than making incorrect ethical judgments.

To see why one might have this worry, we need briefly to consider the nature of disagreement. The sorts of disagreement that will be of interest to us in this context are those that are *substantive*, rather than merely *notional*. As Bernard Williams puts it, for a disagreement to be substantive we must be able to locate a "locus of exclusivity", namely something that one party to the disagreement asserts and the other one denies.[33]

---

[33]See "The Truth in Relativism", in Williams (1981), p.135.

Only if there is a locus of exclusivity can we distinguish cases of 'same-meaning-different-belief' and 'different-meaning' (to use S.L. Hurley's phraseology).[34] To give a very crude and extreme example, if two parties are disputing whether banks contain money the dispute might involve a notional or substantive disagreement. If one of the parties means by 'bank' a financial institution and the other means a piece of land abutting a river, then the dispute will be entirely notional because they are using totally different concepts. On the other hand, the dispute would involve a substantive disagreement if both were talking about financial institutions, but, because of a financial crisis, it was a live issue whether or not the banks were empty. There will only be a substantive disagreement if it is a case of same-meaning-different-beliefs. The possibility that someone is actually mistaken in their use of a concept only pertains if we can distinguish between same-meaning-different-belief cases and different-meaning cases. If we cannot make out this distinction, then anyone applying a concept in a different way could not be regarded as going wrong in their use of the concept, but would be seen merely as using a different concept.

In the context of our discussion, the worry will arise that on a non-reductionist, response-dependent account of ethical objectivity, the requisite distinction between notional and substantive disagreements is not available, since all disagreements are essentially notional. The problem might seem to be that the Internal Criticism Thesis, combined with the Uncodifiability Thesis, ensures that any persistent disagreements must be regarded as purely notional. In the absence of a response-independent standard of correct application of ethical concepts, in cases of sufficiently great disagreement, those at odds with the virtuous cannot be regarded as deploying ethical concepts in at least a

---

[34]See Hurley (1989), Ch.3, and also (1985).

certain range of cases. The virtuous cannot be regarded as being in substantive disagreement with them, because they lack something about which to disagree. If this picture is accurate, then it hardly seems to offer anything worthy of the label 'objectivity'.

McDowell would doubtless respond that this objection caricatures his position, because he accepts that ethics is intelligible to those who are not completely ethically competent (the vast majority of us, no doubt):

> But it is also true, I think, that one does not need the *because* in order to shape one's life as one should; if one's grasp on the *that* is correct, and one acts on it, one will be living in accordance with virtue... [Aristotle] leaves room for a transition to a comprehending acceptance of a scheme of values, and thus connects himself to a tradition that stands precisely opposed to dogmatism.[35]

McDowell's aim is to avert the danger that the account that he takes from Aristotle involves a smugly dogmatic acceptance of the views of a particular group (the virtuous). He hopes to achieve this goal by pointing to the fact that one does not need to understand the *because* that makes an action choiceworthy (something that the virtuous understand), one merely has to understand *that* it is choiceworthy. Both objections that I have raised will fail, McDowell would claim, because it is not just the judgments of a particular group that determine the correct application of ethical concepts, but those judgments that withstand a process of reflective criticism. This criticism will be of the sort illustrated by Neurath's image of rebuilding a boat at sea:

> One reflects on one's inherited scheme of values, or the perceptions of choiceworthiness in action in which the scheme of values expresses itself, from inside the ethical way of thinking that one finds oneself with, not by contemplating it from the external standpoint of a theory about the motivations built into human beings as such.[36]

---

[35]McDowell (1995a), p.212.

[36]McDowell (1995a), p.213.

McDowell could then claim that the judgments that determine the correct application of ethical concepts may be ones that never have been, nor ever will be, made by any real agents.

The Which Ethics? Objection can be rebutted because the ethical way of life is large and non-exclusive. Provided that one has some sense of what is going on ethically, one's judgments are susceptible to Neurathian reflection and are thus capable of moving towards the ideal of the virtuous agent. Those who have judgments that differ from those of the virtuous will be close enough to the virtuous for their judgments to be revisable by Neurathian reflection to bring them towards the moral truth. Those whose judgments are, in principle, incapable of this kind of improvement will be too distant from the virtuous for their judgments to be coherently regarded as ethical in character. The possibility of their existence will, thus, not impugn the objective warrant of the judgments of the virtuous. The Substantive Disagreement Objection fails because even not-truly-virtuous agents share with each other and with the virtuous enough of a common background of judgments that such-and-such is choiceworthy to provide the necessary locus of disagreement. There is thus content to the claim that there is genuine correctness in ethical judgment, since we have a coherent distinction between cases of deploying different concepts and deploying the same concepts incorrectly.

This response certainly has some plausibility. However, I believe that the objections can ultimately both still be pressed by pointing out that reflection aimed at achieving a Neurathian system of coherent beliefs or judgments is itself a process that carries a heavy normative baggage. The Neurathian idea attracts McDowell because it suggests that the form that ethical reflection takes is *sui generis* and does not have to match up to external norms. However, where this process of reflection ends up depends

crucially on what standards one has for coherence, or how one's perceptions "hang together".[37] Different conceptions of what constitutes a coherent and unified set of judgments give rise to very different answers as to what is the objectively correct first-order judgment. The problem now is that the Which Ethics? Objection re-emerges at the next level up, because we can now ask, Which Coherent Set of Judgments?

Different participants in the ethical way of life can have different perceptions of choiceworthiness, which are in need of critical reflection. This fact suggests that the situation will be even more problematic at the more abstract and complex level of considering what constitutes a set of judgments that 'hang together'. Far from resulting in a process of warranted convergence on the truth, the process of reflection that McDowell canvasses could equally result in divergence. Again we seem faced with the dilemma, which McDowell tried to avoid by appealing to Neurathian reflection, of having to choose between dogmatism (the smug assertion that the judgments of a particular group - the virtuous - just *are* the right ones) and the abandonment of the aspiration for objectivity.

A response on McDowell's behalf might be offered to the effect that my objection underestimates the resources available to a coherentist approach and makes the illegitimate demand that one be able to plot out in advance how the process of moral reflection will go. In response to the first point, it is important to note that mine is not an objection to coherentism *per se* - that coherence lacks any bite as a means of delivering objectivity. Rather, the point is that McDowell's Uncodifiability Thesis leaves him in a peculiarly weak position when it comes to spelling out what the *right* sort of coherence

---

[37]McDowell (1995a), p.213.

is that would deliver objectivity. The circularity involved in saying that the right sort of coherence is the sort of coherence that the virtuous would achieve in their reflection is simply too tight to yield anything other than vacuity or dogmatism.

This leads on to a response to the second point. The objection is not that we cannot see how the process of reflection towards coherence is to go - to demand that would be to demand something very like an illegitimate form of codification. Nor is it a demand that we have a conception of right conduct for every situation. I do not want to preclude the possibility that there will be genuine vagueness in how ethical concepts are to be applied or that there are genuine situations of moral conflict. Rather, my question is whether McDowell's account allows determinate answers even in absolutely central cases that cannot plausibly be regarded as lying on a vague or indeterminate periphery.[38] The point of my objection is to question whether the Uncodifiability Thesis allows us to offer any satisfactory account of what it would be for there to be objectivity-sustaining coherence. To return to an analogy I used earlier, the Uncodifiability Thesis might seem to suggest that our understanding of what constitutes success in cricket must be achieved primarily by focussing on the responses of expert cricketers. But this gets things the wrong way round. Success is defined by the rules of the game, and expertise is exhibited in playing in such a way as best to achieve success.

My suspicion is that the Uncodifiability Thesis, read in the strong way that McDowell needs in order to support his claims about the role of virtue, can result in a similarly problematic picture of ethical judgment. The account seems to face an unavoidable dilemma if we take seriously the conceptual priority of virtue, as supported

---

[38]I am grateful to Peter Railton for reminding me of this.

by a strong Uncodifiability Thesis. On the one hand, if we want a reasonably determinate sense of what correct ethical judgment is, then it will seem necessary to demand that the virtuous be a fairly specific, elite group. With the Uncodifiability Thesis in play, this seems to leave us with all the problems of dogmatism and failure to address the Which Ethics? and Substantive Disagreement Objections. McDowell does not seem to want to take this option. Rather, he envisages a wider pool of ethical agents reflecting on what an ideal character would be. This seems more plausible, but the problem is that it risks vacuity. If we are still to take seriously the Uncodifiability Thesis, then it is hard to see how the theory even offers an account of how there might be objectivity in ethics - we seem to be giving the responses of the cricketers the central role in defining what success is in cricket, but with no idea who the cricketers are! This argument does not depend on restricting the reflective resources merely to coherence, for the situation seems even worse if we take seriously McDowell's suggestion that subjects are guided, not just by standards of coherence, but by *logos* or reason.[39] *Logos* is seen as pronouncing upon the good standing of judgments, "using whatever standards it has to hand".[40]

If it were not for the Uncodifiability Thesis, this suggestion might seem helpful. An appeal to *logos* in general (whatever that is) would seem to offer the hope of the sort of absolute vindicatory explanation canvassed in Chapter 4.5. We would have got away from the potentially problematic suggestion that ethics is thought of as a discrete practice that is insulated from other rational practices. *Logos* would be seen as making demands on all our judgments of whatever kind, and virtuous ethical judgment would be seen as a

---

[39]Bill Pollard suggested to me that McDowell should properly be read as offering more than coherence. As I go on to argue, I am not at all sure that this extra does much good.

[40]McDowell (1995c), p.169.

particular instance of the role of *logos* in guiding our rational lives. But it is increasingly hard to see how one could still hold on to a strong version of the Uncodifiability Thesis. This is, in the first place, because correct ethical judgment is being claimed to be responsive to *logos* in general, and not just a specifically ethical conception of virtue. Secondly, the problem of vacuity seems all the worse at this more generalised level. The concept of a virtuous agent seems to have dropped out of the picture, and it seems to be the notion of the demands of *logos* that is doing all the work. This plausible broadening of the conception of the virtuous seems to require a weakening of the Uncodifiability Thesis to something like the Weak Uncodifiability Thesis that I have suggested. But, as we have seen, it is quite compatible with this to envisage the central role in practical rationality being occupied, not by the virtues, but by norms of reason.

I shall end this section by offering a diagnosis of what has gone wrong. The Internal Criticism Thesis made the seemingly laudable claim that objectivity does not require ethical judgment to have any radically external validation. The problems I have been discussing arise when the Uncodifiability Thesis is added, since it either results in an unattractive dogmatism, or else weakens the constraint offered by internal criticism to the point of uselessness. For now it is not only up to ethics to determine what is ethically correct, but it is also up to ethics to determine what would count as an appropriate and successful way of reasoning towards this goal. With the Uncodifiability Thesis in play, only the judgments of virtuous agents (precisely what are to be identified) are available to do this. At this point it is more plausible to think that it is really the Weak Uncodifiability Thesis that is at work, but this offers no distinctive conclusions about the role of virtue in substantive ethical deliberation. One can accept the Internal Criticism Thesis and the basic Neurathian point, but hold that particular normative practices are

answerable to generally-applicable norms governing what counts as appropriate internal criticism.[41] It is possible, for example, that negative recursive principles of rational criticism might play just this role of providing a genuine constraint on reflection about particular judgments. But, as we have seen, this is not a resource that is supposed to be available if one accepts the strong version of the Uncodifiability Thesis.

5. McDowell's Account of Practical Deliberation

My aim, so far in this chapter, has been to question the metaphysical aspect of McDowell's ethical theory, challenging whether the metaphysical ideas he takes from Wittgenstein can really have the implications for substantive ethical deliberation that he suggests. I also questioned whether his account really offers a coherent form of ethical objectivism. In the last two sections of the chapter, I want to continue to pursue the issue of whether McDowell's ethical theory really delivers a genuine form of objectivity, but from a slightly different angle. As we have seen, a response-dependent form of ethical objectivism needs at its heart a theory of practical rationality that can determine which are the relevant extension-fixing responses for ethical concepts. On McDowell's theory, to put matters rather crudely, it is the judgments of virtuous agents that are the relevant extension-fixing responses. It is, of course, also McDowell's contention that an

---

[41]Or, perhaps more appropriately, one might forget the idea that there are separate normative practices in our rational lives and think of reason as one big practice with particular kinds of judgments being normatively responsible to different kinds of reasons. We could then stop thinking of a rational practice as something to which there could be alternatives, and hence drop the idea that there is only one big rational practice, because no sense can be attached to there being an outside to it. See "On the Very Idea of a Conceptual Scheme", in Davidson (1984), and Chapter 6.4 below.

understanding of the demands of morality has to come primarily from reflection on virtuous character.

It will be my contention, in the next section, that the strong Uncodifiability Thesis leads to a conception of practical deliberation by the virtuous that is simply too distant an ideal to offer an account of correct ethical reasoning in the decision situations in which ordinary moral agents find themselves. It thus becomes increasingly problematic to see how agents could structure their deliberation by reflection upon the ideal of the truly virtuous character. Without a viable account of practical deliberation, McDowell's moral realism seems to exhibit what is a serious lacuna for a response-dependent objectivism. We will, thus, have found another way in which the Uncodifiability Thesis undermines objectivist aspirations. We will also have opened up the discussion of what form an objectivist theory of practical deliberation should take, a topic that will be further pursued in the next chapter as I develop my own positive account. This section will be taken up with an exposition of McDowell's account of proper practical deliberation, and it will be in the next that I move on to present my case against it.

Non-cognitivists will claim that the task of providing a cognitivist account of the reasons that agents have for acting is foiled by the fact that it is seemingly perfectly intelligible for two people to be equally aware of the ethical properties of a situation, and yet for one person to be motivated to act in some way, and the other not. They will claim that the only way in which to understand this situation is to regard the two parties as differing with respect to the non-cognitive desires, pro-attitudes or emotive commitments that they have.

The cognitivist or objectivist needs to come up with some other account of this situation. They can begin this task by making two moves. The first is to insist on a

refinement of the role of desires in the explanation of action. The fact that one attributes a desire, as well as a conception of the situation, to an agent, when explaining the reasons that he or she has for acting, does not, in itself, constitute a block to a cognitivist account of practical reason. It does so only if the desire is intelligible independently of the agent's (purely cognitive) conception of the situation, and *vice versa*. For then the conception of the situation can be perfectly objective, but not give the agent an objective reason to act. The conception or belief would form a reason for action only when conjoined with one of the agent's desires. A particular sort of independent intelligibility must be possessed by the desire and the conception, if this move is to be made. The desire has to be intelligible independently of any valuational concepts that might appear in a characterization of the agent's conception of the situation. The conception must be attributable to the agent independently of the agent's having any motivation to act in a certain way. If the desire could not be attributed to the agent without also attributing a certain conception of the situation, or if the conception could not be attributed independently of the agent's being motivated to act, then the objectivity exhibited by the conception of the situation could be carried over to the reasons that it gives the agent.

But *are* there any reasons such as this, which, although requiring that an agent have both a desire and a cognitive belief, do not allow the two to be independently attributable? Thomas Nagel has argued that prudential reasons are of just this kind.[42] When someone acts for a prudential reason, it is enough by way of explanation that we attribute to him or her a certain belief about the situation. We can also ascribe a desire of some sort but, as McDowell puts it:

---

[42]See Nagel (1970), pp.29-30.

> [T]he commitment to ascribe such a desire is simply consequential on our taking him to act as he does for the reason we cite; the desire does not function as an independent extra component in a full specification of his reason, hitherto omitted by an understandable ellipsis of the obvious, but strictly necessary in order to show how it is that the reason can motivate him.[43]

The thought is that the desire is not independent, because, if one were not the sort of person to whom the relevant conception could be attributed, then one would also not be the sort of person to whom the desire could be intelligibly attributed. What makes it correct to say that a desire is intelligible independently of a certain cognitive component (or *vice versa*) is precisely that one can conceive of a person's having the desire without having that particular cognitive component. The claim is that, in the prudential case, this is not possible.

The suggestion that McDowell then makes is that it is not just prudential reasons that exhibit this characteristic, but also moral reasons.[44] His justification for this claim is that, in the explanation of virtuous actions, just as in the explanation of prudential actions, it often suffices to refer to the agent's conception of a situation:

> To explain an action we regard as virtuous, we typically formulate a more or less complex characterization of the action's circumstances as we take the agent to have conceived of them.[45]

But the challenge of how to account for cases where two people have the same conception of the situation, but have different reasons to act, has still not been met. This is where the second of the cognitivist's moves comes in.

---

[43]McDowell (1978), p.15.

[44]See Platts (1981), for a discussion of how hard it really is to come up with genuinely independent, and non-question begging, formulations of the sort of desires that the ethical non-cognitivist requires.

[45]McDowell (1978), pp.15-6.

Non-cognitivists or orthodox Humeans concerning practical rationality will claim that the phenomenon of agents' sharing a cognitive conception of a situation while differing in their reasons to act shows that, what one might call, a *conjunctive explanation* of a person's reasons for action is required. The difference in the two parties' reasons to act is supposed to show that the cognitive conception or belief alone will not do the job: in order to explain a person's reason for action, it must be *conjoined* with some other - non-cognitive - element, namely a desire, which *can* do the job.[46] This move is reminiscent of the way in which the argument from illusion is meant to motivate a certain view of experience in epistemology. On this view, a subject can have exactly the same experience when they are hallucinating as when they are having a veridical experience. It is therefore claimed that there must be some additional element, beyond the experience that distinguishes hallucination or illusion from veridical experience. A popular candidate for this extra element, which, in the case of veridical experience, is conjoined with the neutral experience, is a causal link (of the right sort) to objects in the external world.[47]

McDowell can be seen as offering structurally similar responses to both of these conjunctive accounts. He denies the first part of each of these arguments, namely that in each case there is some common element shared between, respectively, virtuous and non-virtuous actions, and veridical and non-veridical experiences.[48] Instead, the account that

---

[46]It must, *ex hypothesi*, be a *non*-cognitive element or else the two parties would not share the same, purely cognitive, conception.

[47]See Dancy (1995), for a discussion of the way in which a variety of different philosophical positions can be regarded as applications of this argumental form.

[48]For his discussion of the case of practical reason, see McDowell (1978) and (1982b); and for the case of experience, see McDowell (1982a), and also Snowdon (1980-1).

he offers is *disjunctive* in nature. In the case of experience, we should regard the situation of its looking to subject S as if there is a tree as:

either (1) there is a tree making itself manifest to S.

or (2) it is for S as if (1).[49]

In the case of practical reason, McDowell denies that the virtuous and non-virtuous share a common conception of a situation. Virtuous agents see situations as possessing certain ethical features that make certain demands on them, while the non-virtuous fail to act because they see these situations without fully recognizing these ethical features. At worst, they see situations as possessing purely naturalistic properties. At best, they can recognize that there are ethical features to the situation, but may fail to act on them because they fail to see an objective feature of the situation, namely that it makes certain ethical demands on them. A useful analogy would be with someone who looks at a page of mathematical formulae: in the worst case of failure to perceive someone would simply see a page of marks occupying certain spatial positions; better would be someone who saw that the marks were mathematical symbols, and had some grasp of what they meant, but failed to see why one line followed from another.

This is a sketch of the opening moves in one sort of cognitivist attempt to develop a theory of practical reason. McDowell makes these moves, but then develops his theory

---

[49]The fact that there is an analogy between McDowell's treatment of the two issues should not hide the possibility that there may be disanalogies as well. Here is one (suggested to me by Ralph Walker): in the case of action there is no equivalent of something 'being for S *as if* they had a reason to ϕ'. Even if this is a disanalogy, it does not interfere with the point being made. But in fact, far from being a disanalogy, it might be a welcome feature of a disjunctivist cognitivist view of action, since this could well be the appropriate description of a case of akrasia: it *seems* to an akratic that he or she has reached an all-things-considered judgment to ϕ, but the fact that he or she does not ϕ suggests that a full judgment has not been made. The akratic discovers this because he or she ends up not ϕ-ing, whereas someone suffering an illusion may not (but can) be in a position to cross-check with his or her other sense modalities. I shall say more on akrasia in Chapter 6.5(c) below.

in a very particular direction. As we saw in Section 2 above, he has an argument that is meant to show that the issue of good ethical reasoning has to be approached via the notion of a virtuous character. So far we have dealt with the notion of virtue at a rather abstract level, but now we need to explore it in more detail if we are to understand McDowell's conception of practical deliberation. McDowell's theory respects the Socratic tradition of identifying virtue with knowledge of the ethical requirements that a situation generates. Thus, virtue is a state of character "whose possessor arrives at right answers to a certain range of questions about how to behave".[50] Virtue is thus to be seen as a reliable sensitivity to what a situation requires of one. Since McDowell is claiming that there is an identification between virtue and sensitivity to the requirements of a situation, sensitivity to the ethically salient features of situations is not just a necessary condition of virtuous action, but also a sufficient one.[51] Thus, a virtuous agent's reasons for action must be completely explained by the deliverances of the special sensitivities that are the virtues. The virtues have to be identified *en masse* with the special sensitivities to the salient features of situations, otherwise the potential for clashes between the deliverances of these sensitivities would undermine McDowell's claim that virtue issues in right conduct.

So, on McDowell's conception the virtues are special sensitivities to the ethically salient features of situations; they are unified; and the knowledge that they deliver can provide a complete explanation of a virtuous person's reasons for action. But what form does this knowledge take? We might conceive of the explanation of an agent's reasons for action on the model of a 'practical syllogism'. On this model, universal knowledge of what

---

[50]McDowell (1979), p.331.

[51]See McDowell (1978) for a fuller account of how a sensitivity can offer a complete explanation of an agent's reason to act.

one ought to do provides the major premise of the syllogism, and the agent's particular knowledge about the situation provides the minor premise. But, of course, McDowell rejects the idea that our conception of what we ought to do can be codified. As he puts it, "it should seem quite implausible that any reasonably adult moral outlook admits of any such codification".[52] Non-cognitivists will agree that the knowledge that is supposed to supply the major premise for a practical syllogism is uncodifiable, but will see this as an argument against the cognitivist aspiration. They will claim that the 'universal' major premise is not a form of knowledge, but is supplied by a non-cognitive emotive commitment or disposition of the will that an agent possesses. According to their view, one has reason to act only if the purely cognitive and motivationally inert factual conception that one has of one's situation is combined with this emotive, motivationally potent commitment. Of course, we have seen that McDowell rejects as misconceived the idea that rationality demands this sort of codification. He is, thus, able to draw the conclusion that it is not the cognitive nature of the proposed universal premise that is at fault but its supposed codifiability. Thus, the virtues cannot be understood to have their status by reference to some codifiable, reaction-independent general conception of the good. Rather, the virtues must be regarded as sensitivities to the ethical properties of particular situations.

On McDowell's account, correct ethical judgment is the acquisition of an accurate belief (and ensuing motivation to act) through the exercise of a sensitivity to the ethical properties of situations. For virtuous agents this belief will be necessary and sufficient for their acting in the right way. Through the unity of the virtues, it will only be by reference

---

[52]McDowell (1979), p.336.

to the virtuous agent's responses to the ethical features of situations that his or her actions will be explicable. For the virtuous, any non-moral considerations are simply "silenced", as McDowell puts it.[53] In contrast, those who are not truly virtuous will have a rather different structure to their deliberation. Their conception of a situation will be clouded by non-moral factors. Thus, the person who is merely continent will feel the pull of non-moral considerations, but will have the self-control to act in the right way. In contrast, the akratic or incontinent person will have a similarly clouded perception of what the situation requires, but will lack the self-control to act correctly.

I hope that this rather brief and crude sketch of McDowell's subtle and ambitious account of ethical deliberation gives some idea of the target of my criticisms in the next section.

6. 'Silencing' and the Virtuous Ideal

The account of practical deliberation that McDowell sketches is impressively direct in the way that it seeks to secure the objectivity of ethical reasons. One of its attractive features is its economy and theoretical modesty. McDowell's focus on the unified, 'silencing' virtues of the virtuous agent means that there is no need for him to give a messy account of how conflicting virtues can be harmonised. Silencing offers a simple, if ambitious, explanation of the connection between ethical judgment and virtuous action. However, the very neatness of the approach is also one of its most serious problems. In this section, I shall resume my attack on the adequacy of McDowell's theory as a response-dependent form of ethical objectivism or realism. But now my aim will be to question whether the

---

[53]See McDowell (1978), pp.26-9, and (1979), p.335.

account of practical deliberation he offers can play the role demanded of it at the heart of the response-dependent objectivist model. It will be my contention that the conception of the virtuous agent that he offers is simply too idealized to sustain the primacy of virtue in the understanding of ethics, and I shall suggest that the Uncodifiability Thesis is the reason for this problematic idealization.

Let us look again at the central role of the 'silencing' of counter-moral considerations in the deliberation of a virtuous agent, faced with competing considerations about what to do. McDowell regards the virtues as unified. Possession of one virtue goes along with possession of all the other virtues, so there can be no way that the proper deliverances of one virtue clash with the proper deliverances of another.[54] A corollary of this unity of the virtues is the view that deliberation does not involve the weighing of competing reasons, rather:

> If a situation in which virtue imposes a requirement is genuinely conceived as such, according to this view, then considerations which, in the absence of the requirement, would have constituted reasons for acting otherwise are silenced altogether - not overridden - by the requirement.[55]

The idea that the reasons, on which the virtuous person acts, simply *silence* any other considerations to which a situation gives rise creates an image of the virtuous agent that many have found unattractive.[56] It might seem to suggest an insensitivity to genuine features of situations and an ideal of single-mindedness that, at best, pays lip-service to any notion of value pluralism. The suggestion that the ideal of appreciating ethical reasons

---

[54]McDowell (1979), pp.332-3.

[55]McDowell (1978), p.26.

[56]See, for example, Charles (1995), and Wiggins (1995). It is worth remembering that among the attractions of virtue-centred ethics are often cited its supposed sensitivity to pluralism about values and the sensitive treatment it can give to cases of moral conflict.

is an apparently effortless process of perception seems to conflict with the intuition that deliberation is an active process of reflection and weighing of reasons. This is so even if this process of weighing is not to be conceived of in, for example, a crudely utilitarian way.

The objection that the 'silencing' model is unattractive or counter-intuitive arguably has some plausibility, but it is not an objection that I want to pursue. McDowell proposes the virtuous agent as an idealization, and such idealization "is not something to be avoided or apologized for".[57] By their very nature, idealizations will depart from reality, and so the idealization may appear revisionary or counter-intuitive if our understanding of ethics has been fogged by a philosophical preoccupation with moral rules and the weighing of reasons. The challenge to McDowell's account of practical deliberation that I want to offer is a different one, which focuses on the proper role of this ideal.

One way in which an ideally virtuous agent might figure in our ethical practice would be as an example. Virtuous agents might be real people, whose insights and example we might seek out and follow when we are faced with a challenging decision to make. This role might be dubbed 'top-down'- an ethical elite are available to guide the ethical proles. Such an understanding of the role of the virtuous person is not without its problems, but, in any case, it does not seem to be what McDowell has in mind.[58] Rather, McDowell seems to envisage what one might call a 'bottom-up' role for the virtuous person. On this conception, the virtuous agent is only an ideal, not a real person. The strong Uncodifiability Thesis is seen as ruling out the possibility of tackling the question,

---

[57]McDowell (1978), p.28.

[58]One problem, for example, would be how such a person is to be identified by the non-virtuous and, once identified, there might be issues about how they understand the guidance given to them.

"How should one live?", "head on" by giving an answer in terms of moral principles or rules.[59] Thus moral agents must primarily reflect on the sort of character that one should have, the sort of person that one should be. I have, of course, questioned McDowell's arguments in favour of the strong version of the Uncodifiability Thesis; however in what follows I shall assume, for the sake of argument, that it has been successfully justified.[60]

As we saw in Section 4, McDowell is confident that a process of Neurathian reflection by ordinary moral agents can generate an understanding of choiceworthy action, even if they initially lack a fully comprehending grasp of what to do.[61] But what I want to question is how the ideal of the virtuous agent can play any role in this process, given McDowell's understanding of such an agent's idealized deliberation. There is a natural way in which one might understand the proposal that correct conduct is to be understood via the notion of virtue or good character. On this view, moral agents take their conflicting perceptions of the moral demands that a situation places on them and try to develop their sensitivity to these demands in such a way that they are harmonised. There is no guarantee that this process of harmonisation will be entirely successful; indeed given the fact that the sensitivities to different moral aspects of situations may have different origins, it is very unlikely that they will be fully harmonized.

This sort of approach has some plausibility, but it is not clear that it is really available to McDowell. There are two connected reasons for this. The first is that this role for virtue is radically different from the role that virtue plays in the deliberation of the truly virtuous agent. It is the defining feature of virtuous agents, on McDowell's model, that

---

[59]See McDowell (1979), p.347. See also McDowell (1980), (1995a) and (1995c).

[60]See Section 3 above.

[61]McDowell (1995a), p.212.

their virtues are unified, do not give conflicting deliverances and silence any counter-moral considerations. On the more mundane account of reflection on virtue that I outlined in the last paragraph, the possibility of conflict between the virtues seems an essential and unavoidable feature of moral life. This potential for conflict does not appear to be an accidental feature of moral decision, but absolutely central to it.[62] In contrast, McDowell's virtuous agent experiences no such conflict. It is at least tempting to think that McDowell's virtuous agent is an idealization too far and is not even an idealization of *ethical* deliberation at all. McDowell's idealization seems to be somewhat akin to an aeronautical engineer designing an aeroplane under the idealizing assumption that there is no gravity. This might make the design neater, but it does not make for very airworthy aeroplanes.

But maybe this objection does not take seriously enough the fact that the virtuous agent is an ideal and it seeks to elevate to an essential feature of ethical decision-making what is really a feature of our messy, non-virtuous deliberation. This point may or may not impress, but I think that my challenge to McDowell can be developed to deal with this point. Let us accept, for the sake of argument, that the virtuous agent is an ideal and that our attempts to decide on choiceworthy action must fundamentally be an attempt to come to an understanding of what this ideal is, that is, what one of virtuous character would do. The problem now is how this deliberation is to be conducted. As we saw in Section 4, the process of reflecting on how our perceptions of choiceworthiness "hang together" will carry a heavy normative baggage - how is this reflection to be conducted correctly? Now,

---

[62]This should not be seen as advancing the claim that values can be deeply incommensurable. They may or may not be, but it is compatible with values being commensurable that such comparison can be difficult to effect.

if the answer is that it must be conducted in the way that the virtuous agent would reflect, then this is doubly unhelpful. It is unhelpful, first because it does not tell us very much, but secondly, and fatally, because this is precisely *not* the sort of reflection in which, *ex hypothesi,* a virtuous agent engages! Because of the silencing effect of their all-things-considered judgments, virtuous agents do not have to deliberate over the deliverances of competing and conflicting ethical sensitivities.

Because the virtuous agent's ethical deliberation is marked by silencing, this ideal is simply too radically distant from ordinary (and indeed, not so ordinary) moral practice to offer a route to an understanding of how one should live via an understanding of the virtuous person. This rigorous interpretation of the virtuous ideal seems to offer little more guidance than the injunction 'Do as someone who does the right thing does'. If the ideal is so distant, then ordinary moral agents, reflecting on their imperfect and conflicting moral perceptions, might as well give up reflecting about the unhelpful notion of virtue. Instead, they might concentrate on what reasons speak for or against acting in one way rather than another. They might, by this route, end up with an understanding of what a virtuous person would do, or, more ambitiously, what a virtuous life would be, but this would merely be a by-product of their deliberation. It seems to be the Uncodifiability Thesis that is to blame for the failure of McDowell's conception of the role of the virtuous agent. First, it is supposed to rule out the possibility of tackling ethical questions other than via an understanding of the ideal of the virtuous agent. Then, by allowing us to say little about the deliberation of the virtuous save that counter-moral considerations are silenced, it radically, and fatally, removes that ideal from us. The result would seem to be that McDowell ultimately has an empty hole where a theory of practical rationality would need to be in a viable response-dependent objectivist theory.

It might be argued, at this point, that it does not matter from the point of view of developing a *response-dependent* objectivism if the concept-fixing responses are those of idealized subjects. This is may be true, but not very helpful. First, we can be suspicious of a response-dependent objectivism that gives a central role to the responses of radically idealized subjects. When the subjects are so idealized (as I have argued, the virtuous are on McDowell's account) we can say little more about their responses than that they are the responses that a perfectly correctly reasoning subject would give. There then seems to be little to be gained by adverting to the virtuous rather than merely saying that they are the responses that practical rationality demands. This is very much the point that I made in the previous paragraph. Secondly, even if such a response-dependent objectivism were viable, it would still be problematic how we are to reach our understanding of the demands of morality via "the notion of a virtuous person", as McDowell hopes.

However, there is a possibility that I have completely misrepresented the role of silencing in McDowell's account of deliberation by the virtuous. So far I have been regarding it as what deliberation by the virtuous consists in - that is, when the virtuous deliberate, their unified virtues simply silence certain considerations. Unlike the ethically less competent, who weigh competing considerations, reflect and, perhaps, engage in a spot of universalization, the virtuous person simply perceives which course of action is choiceworthy. Here silencing features as a central part of the process of deliberation leading up to action. However, one could regard silencing as a feature of the relationship between a virtuous agent's all-things-considered judgment as to what to do and his or her motivation to act. For the virtuous person, the all-things-considered judgment, once made,

simply silences, and renders motivationally inert, any competing considerations.[63] On this view, silencing would be detachable from McDowell's favoured perceptual metaphor. McDowell does not give an account of how the virtuous person's practical deliberation proceeds, but any considerations that compete with the output of this deliberation (whatever it is) are silenced.

Ultimately, I do not think that this sort of reading makes the theory any less problematic. There are two reasons for this. First, if we conceive of silencing as a feature of the relationship between virtuous agents' all-things-considered judgments and their motivation to act, then it ceases to be a feature that is unique to a neo-Aristotelian virtue-based conception of deliberation. A consequentialist or utilitarian moral theorist would claim that a virtuous person is one who always acts to maximize utility. They can then say that, for the virtuous, all other, non-utilitarian considerations are simply silenced by the all-things-considered judgment that a certain course of action is utility-maximizing. In fact, it seems to be very easy to see how such 'silencing' might take place on such a theory (easier, arguably, than on a virtue theory). If the only thing that matters is that utility is maximized, then this is all that a virtuous agent needs to be sensitive to. In different circumstances, features of a situation that make a certain course of action utility-maximizing would cease to do so and hence would be morally 'silent' as far as a different all-things-considered judgment about choiceworthiness is concerned. This would be a rather surprising result. McDowell seems to think that silencing is a phenomenon that is special to his brand of Aristotelianism. However, if we treat silencing as a feature of the

---

[63]I am grateful to Gerald Lang for pressing this objection.

relationship between judgment and motivation, then even McDowell's arch-opponent, the utilitarian, could find room for the phenomenon.

The second problem with this way of understanding the role of silencing is that it seems to put the process of practical deliberation by the virtuous into a 'black box'. We know something about the features of the deliverances of this process, namely that they silence any counter-moral considerations, but we are still lacking any account of the workings of the process of deliberation itself. This would appear to leave us with an even worse problem than when we understood silencing as what virtuous deliberation is. On our new reading of McDowell, he seems to say nothing at all about the process of ethical deliberation by the virtuous. Again, the reason for this deafening silence, is, presumably, that the strong Uncodifiability Thesis rules out saying anything very informative about the deliberation of the virtuous, except that it will yield the correct answer. To the sceptic it is going to look suspiciously as though McDowell's theory is so modest that its promise of providing a theory of ethical deliberation is all but empty. One has to be very confident of the grounds for the strong reading of the Uncodifiability Thesis if one is happily to accept this extreme theoretical modesty.

Given that we are to understand ethics via "the notion of a virtuous person", either of our two ways of understanding the role of silencing seems to leave us with a worrying lacuna in the response-dependent objectivist model.[64] We seem either to be left with no theory of practical rationality to drive a response-dependent objectivism, or else one that is so idealized that it cannot provide a route to answering the question, "How should I act?". This second option might, possibly, allow a response-dependent objectivism to get

---

[64]The quotation is from McDowell (1979), p.331.

off the ground, but only at the cost of making the role of the virtuous ideal radically self-effacing. So, not only have we found reasons to doubt the plausibility of McDowell's conception of the role of virtuous deliberation, we have also discovered another way in which its objectivist aspirations come to grief. Again, one of the key contributors to these woes is the Uncodifiability Thesis.

7. Conclusions

This has been a long chapter, and the argument has been advanced on several different fronts, so I shall attempt to draw the strands together. I have tried to shed some light on what might be involved in developing a non-reductionist, response-dependent conception of ethical objectivity in a virtue-centred, neo-Aristotelian manner. We have looked at how this approach attempts to give an account of the objectivity of ethical judgment and at how it offers a theory of practical deliberation, and I have found it wanting in both regards. In particular, four tentative lessons can be learnt:

(a) McDowell tries to draw a lesson from the nature of rules as they appear in the *recherché* context of Wittgenstein's rule-following considerations and then tries to use this to make a distinctive claim about how first-order ethical deliberation should be understood. In contrast, I have suggested that there is no need to think of ethical rules or principles as part of some metaphysically dubious, external validation of ethical practice. Constitutive rules seem to be entirely different beasts from those that are the target of Wittgenstein's rule-following considerations. Similarly, recursive negative constraints on rational criticism are appealed to precisely when one has rejected any kind of foundationalist, external grounding for reason. The case for choosing McDowell's

preferred, highly modest form of neo-Aristotelianism, rather than something a bit more conceptually ambitious, does not seem to have been provided by Wittgenstein's rule-following considerations.

(b) The Uncodifiability Thesis restricts the conceptual resources available for internal criticism to such a point that it is hard to sustain a distinction between judging wrongly and simply judging using different concepts. Without this distinction in play, there can be no real sense in which subjects can be objectively correct or incorrect. Thus, ethical objectivity does not appear to be achievable under this constraint. McDowell's hope that the ability of virtuous agents to engage in Neurathian reflection will plug this gap might seem to be over-optimistic. With the Uncodifiability Thesis in place, what agents come to regard as a properly coherent set of judgments or beliefs will itself be dependent solely on the particular normative practices to which those agents belong. There can be little confidence that the fragmentation in judgment that occurs at the first-order level will not extend to, or even be magnified at, this more abstract level.

(c) The apparent incompatibility of the Uncodifiability Thesis with a genuinely objective, response-dependent view of ethics suggests that there needs to be some sort of conceptual gap between the judgments of agents (even of ideal agents) and the standard of objective correctness in judgment. The provision of a more theoretically articulated theory of practical rationality at the heart of the response-dependent objectivist model seems a possible way of satisfying this desideratum. In the next chapter, I shall explore the idea that a solution of this kind can be developed by appealing to recursive norms of rational criticism.

(d) I have argued that there are grounds for thinking that we should accept McDowell's insight that, in explaining why a person acts morally it can be enough simply to cite the purely cognitive conception of his or her situation that made an action choiceworthy.

However, I have suggested that McDowell's account of the virtuous person's conception of a situation and how this ideal is related to our actual reason-guided deliberation is deeply flawed. The Uncodifiability Thesis seems to rule out giving any account of how the virtuous deliberate practically, beyond describing certain features of the output of that deliberation. But, understood in this way, it is hard to see how the ideal of virtuous character can structure the practical reflection of ordinary ethical agents, as McDowell thinks it does. Given the austerity of McDowell's account, the virtuous ideal seems simply too radically distant from the deliberative situation of ordinary ethical agents. In the next chapter, I shall suggest as an alternative a model that is cognitivist, but in which a process of rational reflective criticism guided by negative regulative norms is absolutely central.

Chapter 6

TOWARDS A 'KANTIAN' ETHICAL COGNITIVISM

1. Taking Another Approach to Practical Rationality

We have now looked at one way in which a response-dependent conception of ethical objectivity might be developed within a broadly Wittgensteinian framework, and we have found it to be quite seriously wanting. In this chapter, my aim is to propose an alternative approach to the account of practical rationality that lies at the heart of the response-dependent model. A key lesson of the previous chapter was that, to the extent that Wittgenstein's rule-following considerations do support the uncodifiability of rationality, it is not an uncodifiability that need be problematic for the theoretical centrality of principles or norms of reason. In Section 5.3(b), I introduced, as a variety of principle that was immune to problems of uncodifiability, a conception of negative regulative principles of reason. In this chapter, I shall pursue this possibility and explore how Kant's Categorical Imperative (CI) might be treated as a regulative principle of reason that can play a fundamental role in the required theory of practical rationality.

The introduction of the CI might seem to be a quixotic and arbitrary development, for it seems a long way from the sort of virtue-based conception of practical rationality that I have been discussing. In fact, I hope to show that it is a surprisingly natural, and theoretically economical, move. The virtue-centred theory of practical rationality offered a unified account of the objective correctness both of agents' reasons to act and of their ethical understanding of their situation. Correct judgment about what to do takes the form of a true belief about what a situation demands. Correct judgment about the valuational

features of a situation also takes the form of true belief. We found, however, that, ultimately, the strong interpretation of uncodifiability, which was meant to support the modesty of the virtue-centred approach, also undermined its objectivist aspirations. Kant's CI is often seen as the core of a deontological, duty-based morality that is radically opposed to the particularist virtue-centred approach. However, the CI is best regarded, not as the generator of inflexible and absolute duties, but as a regulative norm for the reflective criticism of one's reasons. Viewed in this light, the CI seems entirely compatible with the level of uncodifiability of rationality that the rule-following considerations seem genuinely to suggest. It might offer, I shall argue, exactly the sort of criterion of correct judgment that a reliance solely on the sensitivities of idealized virtuous agents was ultimately unable to provide. This Kantian approach can, thus, be seen to offer a rather more satisfactory account of the theory of practical rationality at the heart of a response-dependent ethical objectivism. In many ways, it is also a natural way of dealing with the inadequacies of a virtue-centred approach, whilst respecting a non-reductionist conception of rationality that eschews strong codifiability.

The CI has traditionally been regarded as a method for assessing the moral acceptability of agents' maxims, or reasons, for action. However, I shall argue that it can also serve as a criterion of correct judgment for agents' appreciation of the ethically salient features of their situation. Serving this role, it can sustain a level of objectivity in agents' ethical beliefs that warrants the label 'cognitivist' for the resulting model. In the next section, I shall introduce the CI as one way of developing the idea that objectivity in judgment requires that there be nothing else to think but a certain (and the correct)

answer.[1] I shall then go on, in Section 3, to look at how the CI can be developed as a criterion for critical reflection, on not just agents' reasons for action, but also on their ethical beliefs. Because it can give an account of correct ethical belief, we are able, with some justice, to regard the account as a form of 'Kantian' cognitivism. I then go on, in Section 4, to look at how this Kantian cogntivism fares in meeting the original challenge of delivering an account of the way in which ethics could admit of objective judgments. Section 5 contains some suggestions about how this broadly cognitivist account of practical rationality might offer potentially fruitful ways of addressing some issues in moral psychology and meta-ethics that have proved troublesome for objectivist theories. Finally, in Section 6, I attempt to draw together the conclusions of the entire thesis and highlight the potential advantages of my proposal, as well as areas that could, and should, be explored in more detail.

There will, I am sure, be other ways in which we could move on from the point that we have reached in the argument, since there will be other options to pursue in developing the requisite theory of practical rationality. One might, for example, try to develop a different sort of virtue-centred account, perhaps one that relied upon a reaction-independent conception of the good or of human flourishing. Or some form of utilitarianism might be thought to play the desired role. It would take me too far off route to try to flesh out, and then rebut, these alternative approaches in any detail. Others are perfectly at liberty to pursue these options within the theoretical constraints that we seem to have unearthed so far. I shall pursue the Kantian approach, because it seems to offer the theoretical benefit of requiring very modest extra conceptual resources and because

---

[1]See Chapters 2.7 and 4.4-4.5 above.

there also seem to be reasons to suspect that the alternatives will be inferior in important respects. I shall, very briefly, highlight these reasons.

A virtue-centred theory might seem to require an even smaller additional theoretical step to be taken than does the approach that I am proposing. Whilst this might be true of a virtue-centred theory that is response-dependent, I suspect that much of the force of the arguments that I deployed against McDowell's account will also apply to other versions of this approach. If the virtue-centred theory is of a response-independent variety, then this will certainly require a considerably larger additional conceptual step to be taken, thus losing out on any claims to theoretical economy. It is, also, not obvious that a response-independent conception of human flourishing would be compatible with the level of uncodifiability that we have found to be genuinely suggested by Wittgenstein's arguments. It is also open to question how plausible a response-independent conception of human flourishing will be when it is shorn, as it surely must be, of the sort of teleological biology that Aristotle envisaged.

A broadly utilitarian conception of practical rationality also appears to be a potentially problematic way of advancing along the response-dependent objectivist route. Utilitarianism is usually associated with a non-cognitivist meta-ethics and has a strong tendency to be reductionist in nature. It is possible to have cognitivist versions of utilitarianism, but the theory's usual reductionist tendencies would seem to be difficult to reconcile with a strongly non-reductionist Wittgensteinian approach to objectivity and rationality. In particular, it will be harder for a utilitarian theory of practical rationality to avoid exhibiting an illicit form of codification. Both a virtue-centred account of practical rationality, whether response-dependent or response-independent, and my proposed Kantian account aspire to give a unified explanation of how practical rationality can play

a role in determining both correct reasons for action and the appropriate understanding of the valuational features of situations. This promises a pleasing theoretical neatness for an objectivist or cognitivist account of practical reason. However, it is not so obvious how a utilitarian account could offer such a unified explanation of both the axiological and reason-giving aspects of practical rationality. While there might be some plausibility in the proposal that the morally right course of action is the one that is utility-maximising, a utilitarian theory does not seem likely to be able to offer an account of accurate ethical situational appreciation. Such a theory will either need to borrow its account of this from elsewhere, or else (perhaps for reductionist reasons) regard the recognition of ethically salient features as essentially unproblematic.

These remarks are extremely sketchy and dogmatic, but they do offer some indication of my negative reasons for preferring my chosen approach over other obvious candidates, reasons that I would pursue if I had the space. In any case, my aim is to sketch an account of practical rationality, and to highlight its advantages and, also, areas where it might run into difficulties. It is not my aim to demonstrate conclusively its superiority over potential rivals.

As, I suspect, this preamble makes clear, the arguments in this chapter will be considerably more speculative and exploratory than those in previous chapters. I hope that what these arguments lack in rigour is made up for, to some extent, by their suggestiveness and any illumination that they might offer.

2. Objectivity and Kant's Categorical Imperative

In Chapter 4.4, I suggested that a fruitful way of thinking about truth and objectivity in areas where a form of scientific realism is clearly inappropriate is to pursue a proposal of

David Wiggins's. He, of course, suggests that a subject matter admits of objectivity if it is possible to provide vindicatory explanations of why there is nothing else to think but that p, where p is the objectively correct answer.[2] One way of understanding the uncodifiability that McDowell thinks attaches to rationality, and ethics in particular, is to see it as severely restricting any general theoretical account of the form that these vindicatory explanations might take. However, if, as I have suggested we should, we do not take ethical theorising to be restricted by such severe constraints of uncodifiability, what can we say about the form that such vindicatory explanations might take?

The schema "there is nothing else to think but p", can naturally be expanded to include things we might do, other than just thinking. If there is nothing else to think but that torturing cats is wrong, we can say (slightly unnaturally, admittedly) that there is nothing else to *do* but not torture cats. And quite generally it seems that we could make this move from what we have reason to think to what we have reason to do, say, want or, indeed, anything else that might be open to reason and that might figure as part of the proposition p in Wiggins's schema. So, it is natural to widen Wiggins's proposal to one that says that we have an objective reason to do, say, want, or whatever, if we have a vindicatory explanation that shows that there is nothing else to do, say, want, and so on, but p. We thus have a very general schema for examining whether we have objective reasons, both practical and otherwise.

The next step we can take is to see that if there is nothing else to think, want, do, and so on, but p, then there is some sense in which it is impossible to think, want, do, and so on, other than p.[3] Clearly, what this impossibility comes to will be crucially important.

---

[2]See Wiggins (1998), p.153, and (1990-1), p.66.

[3]See Moore (1996), pp.166-9.

Before we address this issue (rather on the hoof), I want to introduce the proposal with which the rest of the chapter will be concerned. In Chapter 5.3(b), I introduced Onora O'Neill's suggestion that Kant's CI is best interpreted as a negative recursive norm of rational criticism, which requires that "any principles of thinking and acting that can have authority cannot enjoin principles on which some members of a plurality cannot (not 'would not'!) act".[4] So, the proposal is that Kant's CI is an obvious way of understanding the idea that objective reasons demand that there be nothing else to think or do.[5] Kant's CI is primarily (and traditionally is seen to be exclusively) a tool for assessing the permissibility of maxims of action. However, the route by which it has been introduced here suggests that it might also be used to test the permissibility of what we think or believe, and hence might be a test for the objectivity of the reasons that we have, not just to act, but to believe certain things. And this might lead us on to a form of ethical cognitivism. We will investigate this possibility in the next section.

In this section, I shall first sketch some reasons why one might understand the CI in the way that O'Neill suggests, and then go on to examine how the CI can yield determinate results as a test of reasons for action. The key challenges that this proposal faces are whether it is possible to develop a plausible conception of the impossibility of thinking or doing other than p, and whether it is possible for its relatively sparse resources to generate determinate answers. If the impossibilities in action and thought can be generated without an appeal to non-objective, unvindicated factors, and if determinate results are yielded, then we have a claim to have found an objectivist conception of ethical

---

[4]O'Neill (1989), p.20.

[5]Moore (1996) makes some intriguing suggestions about how the "nothing else to think" schema might be applied and extended, in part in Kantian directions.

judgment. The proposed approach avoids the traditional charge of empty formalism levelled at Kantian ethics, but introduces problems of its own. The move towards a more cognitivist approach will be seen to defuse these problems somewhat.

It is possible that I have already said enough to make plausible the suggestion that the CI can be seen as one way in which to understand the "nothing else to think" schema.[6] However, this might seem a surprising proposal for those wedded to the idea that the CI is a purely moral device that functions as a generator of duties. So, I shall briefly introduce O'Neill's justification for this understanding of the CI.[7] In O'Neill's view, Kant regarded the critique or vindication of reason as a practical task, governed by the negative regulative maxim of the CI. So the CI is seen as the supreme principle of *all* reason. If reason is to be vindicated, how might this task be pursued? One way might be to approach the vindication in a foundationalist manner, starting with some fundamental principles and working up from these. But a 'vindication' that rests on unvindicated assumptions is no vindication at all. An alternative form of foundationalism would be to vindicate reason by analogy with power or force.[8] This force (if it is to be what vindicates reason in a foundational way) must be alien to reason, but, in that case, the appeal to an alien force or power would, far from vindicating reason, undermine it:

---

[6]It is not necessary, for my purposes, to establish whether or not O'Neill's Kant is the true Kant. What matters is the independent plausibility of the sketched line of thought and its compatibility with the proposed understanding of objectivity encapsulated in Wiggins's schema. I take it that even if this is not the real Kant, it is Kantian enough for my purposes.

[7]See O'Neill (1989), Part I (especially Ch.1), for this argument in more detail.

[8]The thought here is that reason might hold its authority in the manner of a dictator who holds power by means of force. One might see reason's authority as taking this form if one regards it as grounded in essentially unvindicatable traditions. Another aspect of this idea is that reason is merely rhetoric - its authority is gained merely by bullying and bamboozling others into submission.

> Reason must in all its undertakings subject itself to criticism; should it limit freedom of criticism by any prohibitions, it must harm itself, drawing upon itself a damaging suspicion... Reason depends on this freedom for its existence. For reason has no dictatorial authority; its verdict is always simply the agreement of free citizens, of whom each one must be permitted to express, without let or hindrance, his objection or his veto.[9]

On this view, reason's authority is seen as analogous to the authority of a decision reached by a plurality of persons relying on no previously fixed harmony of opinion, and any appeal to alien authority only weakens the authority of reason.

If reason is not to be vindicated foundationally, then another possibility is the sceptical one that no vindication of reason is possible. But Kant found such scepticism intolerable.[10] So, if a vindication of reason cannot rely on the alien and unvindicated, then, O'Neill suggests, by a process of elimination we have to see the authority of reason as self-imposed and its vindication as a recursive process, relying on no predetermined (and hence unvindicated) procedures or algorithms.[11] One might still wonder why, even if the vindication has to operate on pre-existing materials, the method by which the vindication proceeds must be so strongly recursive that the method itself is open to revision. Why, one might wonder, cannot appeal be made to algorithms that stand outside the scope of such revision? But, according to O'Neill's Kant, principles of reason cannot be entirely algorithmic, because algorithms are incomplete and require premises that must either be unvindicated or else be supplied by reason.[12]

So, the conclusion is that principles of reason must be negative requirements that leave matters largely indeterminate. The requirement that reason should not rest on alien

---

[9]Kant (1929), A738/B766.

[10]See Kant (1929), Avii, for an expression of the intolerability of scepticism..

[11]See O'Neill (1989), pp.22-3.

[12]See O'Neill (1989), p.19.

authority or force is just such a negative constraint. O'Neill suggests that Kant's use of political metaphors in his discussion - metaphors such as those of tribunals, debates and communities - is vital to our understanding of how the reflexive vindication of reason is to be conducted.[13] He sees the project of the critique as being conducted by a plurality of "citizens" and "fellow workers".[14] Since there can be no prior determination of the positive method by which this plurality engages in the critique, the metaphors of a tribunal and a debate are particularly apt: a tribunal's decision is left largely under-determined by any procedural rules that might cover its operations; and, similarly, a debate is to be regarded as free and open-ended, with questions of correct procedure themselves falling within the scope of debate. No decision is utterly fixed since the recursive nature of a free and open-ended debate means that decisions are always open to revision in the light of further debate.

In this context, O'Neill sees the CI - the principle that no principle of thought or action should be accepted that *cannot* be accepted by every other member of the plurality - as another negative requirement on reasoning.[15] Without such a principle, communication and interaction between beings who are not guaranteed by inclination to agree would be impossible.[16] Such a principle is a purely negative requirement on any possible debate between a plurality of non-coordinated subjects, and, as such, it leaves

---

[13]See O'Neill (1989), pp.17ff.

[14]See O'Neill (1989), p.22.

[15]See O'Neill (1989), p.20-1.

[16]My discussion of *PI* 242 in Chapter 3.6 suggests that this need not conflict with Wittgenstein's claim that there must be *agreement* in judgments, because the sort of agreement for which Kant says there is no prior guarantee is agreement in opinions - for which reasons can and must be given. This would seem to be something that, from a Wittgensteinian perspective, one could accept.

largely under-determined what the actual outcome of any debate will be. Despite this under-determination, the CI nonetheless has bite, excluding as it does, modes of thought and action that are in principle unshareable.[17]

With a slightly better idea of the rationale behind the proposed understanding of the CI, I now want to turn to the question of how a mechanism that seems so sparsely resourced can deliver determinate answers about the permissibility of our actions.[18] According to Kant, when any rational agent acts intentionally he or she acts on a maxim[19]. A maxim is not just any old intention that an agent has, but, as O'Neill puts it, "those underlying principles or intentions by which we guide and control our more specific intentions".[20] Thus, in writing this chapter, I might be acting on a maxim of writing a penetrating and incisive discussion of some key issues in Kantian ethics. This will guide me in the formation of other, subsidiary intentions concerning which articles or books I choose to read, which specific topics to address as I write, and when to decide that I have read and thought enough and that it is time to switch on my laptop. This distinction

---

[17]All this talk of vindicating reason might set alarm bells ringing, given the broadly Wittgensteinian conception of objectivity that I have developed. Although I cannot properly defend the claim, this fear seems misplaced. This explanation of the role of the CI is strongly non-foundationalist and firmly eschews the codifiability of rationality, and, thus, should be highly congenial to this wider project. There may, however, be a concern that Wittgenstein stressed the key role of unreflective action and that this is incompatible with the project of vindicating reason. However, I think that this concern is motivated by a temptation either to paint Wittgenstein as offering a (reductionist or non-reductionist) constructive account of normativity in terms of communal agreement, or as a quietist. In Chapter 3, I sought to reject all of these readings of him. In any case, if an aspiration to vindicate reason is felt to be too grand, we can satisfy ourselves with the fairly obvious connection between the CI (as interpreted) and Wiggins's proposal.

[18]My discussion draws on the conception of maxims and the operation of the CI offered by O'Neill, especially in "Consistency in Action" and "Universal Laws and Ends-in-Themselves" (both in O'Neill (1989)).

[19]Since a maxim is a "subjective principle of action" (Kant (1988), IV, 421), and one must have a subjective principle in order to have an intention.

[20]O'Neill (1989), p.84.

between maxims and mere intentions is crucial. If there were multiple maxims for any one act (as would be the case if they were basically just ordinary intentions), then the CI could not guide action. This is because there might be some maxims for an act that passed the test and others that did not, leaving the agent without guidance as to what is required. But if a maxim is the supreme guiding principle on which an agent acts, then there will only be one of them for any one intended act, and so the problem of multiple maxims simply does not arise.

Kant's focus on maxims as the primary objects of moral assessment has important implications for the viability of his theory. To establish whether a maxim is a morally permissible one on which to act, Kant claims that it has to be able to pass the CI test, of which he gave three main formulations that he believes are "at bottom merely so many formulations of the same law".[21] I shall not look at the issue of the equivalence, or otherwise, of these formulations and shall, indeed, focus on only one formulation, the Formulation of the Universal Law (FUL):[22]

> Act only on the maxim through which you can at the same time will that it should serve as a universal law.[23]

For a maxim to serve as a universal law, it must be one upon which all rational beings could act. The key question is now, What constraints are there on one's willing a maxim to be a universal law? As it stands, Kant's FUL might look as if it is asking what we would *want* to hold as a universal law, in much the same way as, say, R.M. Hare's

---

[21]Kant (1988), IV, 436.

[22]Although, on the question of the equivalence of the formulations see, "Universal Laws and Ends-in-Themselves", in O'Neill (1989). It would, doubtless, be fruitful for a full treatment of this approach to examine what additional light the Formulation of the End-in-Itself and Formulation of the Kingdom of Ends can add.

[23]Kant (1988), IV, 421.

universal prescriptivism demands that we be able to universalize our non-cognitive prescriptions.[24] The question then irresistibly arises of what account is to be given of rational wanting. If one takes this line, then a serious problem arises for Kant, since he claims, in the first chapter of the *Grundlegung,* that it would subvert the rational nature of morality if it were founded on anything contingent, such as the possession of desires or inclinations.[25]

But this would be quite seriously to misread Kant's position. He is not saying that one should ask a *hypothetical* question - What *would* I will, or be able to will, if I were to consider the wishes of all other parties? This would then require a specification of what powers of willing (construed as wanting) would be available in a certain hypothetical situation. Rather, Kant is asking a *modal* question - What *can* I will?[26] And when we look for a constraint on what one *can* will we are not to look at how 'strong' one's will is. Instead, Kant focuses on what it would be conceptually impossible to will.[27] So it is mistaken to think that some account of one's emotional or volitional strength is needed. Kant proposes two tests to see if the universalized willing of a maxim is rationally impossible - a test for a contradiction in the conception of a maxim, and a test for a contradiction in the willing of a conceptually coherent maxim:

> We must be *able to will* that a maxim of our action should be a universal law - this is the general canon for all moral judgment of action. Some actions are so constituted that their maxim cannot even be *conceived* as a universal law of nature without contradiction, let alone *willed* as what *ought* to become one. In the case of others we do not find this inner

---

[24]See Hare (1963) and (1981).

[25]See Kant (1988), IV, 389.

[26]See, especially, O'Neill (1989), pp.138-9.

[27]The connection with the impossibility in thinking in Wiggins's schema that I am trying to explore should be clear.

> impossibility, but it is still impossible to *will* that their maxim should be raised to the universality of a law of nature, because such a will would contradict itself.[28]

So, some maxims are claimed to be so internally incoherent that we cannot even get to the stage of willing them. Examples of such maxims that contain contradictions in their very conception would be maxims to engage in theft and maxims to borrow money when one is in need, but has no intention of returning it, despite having promised to do so.

Let us examine these examples in turn. Suppose I form a maxim, "To engage in acts of theft when it would be to my advantage to do so", and try to universalize it, producing a universal law of nature, "Everyone will engage in acts of theft when it is to their advantage". Now, theft is the unlawful taking of property. If this universal law were to hold, then the normal and predictable consequence, which the universal legislator (in whose role I am acting when I try to universalize a maxim) would recognise, would be that the institution of property would cease to exist (at least in any recognizable form).[29] So the maxim contradicts itself - theft requires property, but universalized theft destroys the institution of property.

At this point, it might be objected that property is a contingent human institution and, surely, this is the sort of thing that Kant is trying to eschew in his approach to ethics, since he seeks moral laws that hold for *all* rational beings, *qua* rational. But this misses the point that, although the CI holds for all rational beings, the *specific* duties that it dictates will vary between species of rational beings, depending on their needs, abilities, inclinations, and so on. Human beings have a natural desire for their own happiness, therefore there is no need for a principle of duty to seek one's own happiness. Creatures

---

[28]Kant (1988), IV, 424.

[29]See Nell (1975), pp.63-81 (especially pp.70-1).

whose mental states were perfectly transparent to each other would have no need for duties not to lie when they communicate. For omnipotent and totally independent beings, there would be no need for duties of benevolence, whereas for finite and vulnerable rational creatures, such as ourselves, these duties are necessary.[30]

But the argument against theft can be made out without making any direct appeal to the institution of property, without making it look as though the institution was playing a specifically moral role of its own in getting the requisite conclusion from the CI test. The normal intention with which one engages in theft is that one should materially benefit from it (indeed this is explicit in the sample maxim that I formulated). But one would not be able to benefit from theft if there was a universal law that everyone should steal when it suits them, since what one had stolen would then be under threat from further theft. Here, it comes out explicitly that the purpose that is served by my theft, namely my material advancement, also requires an institution, such as that of property, that protects me from others, and the universalization of the maxim would undermine this protection.

A similar story can be told when one considers (in the role of universal legislator) the maxim, "To borrow money when in need and promise to return it but have no intention of doing so".[31] As universal legislator, one considers the law, "Everyone to borrow money when in need and promise to repay even when they know that they will not". One also considers the normal subsidiary intentions that accompany the maxim and the normal consequences that would stem from the proposed law holding for all rational beings. Clearly, the institution of borrowing requires a reciprocal institution of lending,

---

[30]See, for example, Kant (1988), IV, 397, for the suggestion that the good will is exposed to "certain subjective limitations and obstacles". On the finitude of human rational agency, see O'Neill (1989), Ch.7.

[31]Kant (1988), IV, 422.

but such an institution would collapse if no one ever repaid their debts. Lending simply *is* giving to someone else the temporary use of something that you possess, and, if that use becomes permanent, then the activity would be giving, not lending, and this is an activity that would not be so widely practised. Again, as with theft, such a situation would be contrary to the normal intentions with which the maxim was formed. The maxim of not repaying one's debts is typically formed with the intention that one's self-interest be served in this way. However, the collapse of the institution of freely lending to those in need would thwart the achievement of this intended goal for vulnerable, mutually dependent beings, such as ourselves.

Kant's focus on maxims that are the subjective principles of action solves the problem of relevant descriptions, which leads to the charge of empty formalism. We can now see why an objection, such as the following classic one from Hegel, is misplaced:

> [I]f the definition of duty is taken to be the absence of contradiction ... then no transition is possible to the specification of particular duties nor, if some such particular content for acting comes under consideration, is there any criterion in that principle for deciding whether it is or is not a duty. On the contrary, by this means any wrong or immoral line of conduct may be justified ... the absence of property contains in itself just as little contradiction as the non-existence of this or that nation, family, etc., or the death of the whole human race.[32]

Hegel's claim is that, taken on its own, a situation such as the absence of the institution of property is not flatly contradictory, and, without some extra element added to the story, one can imagine a person cheerfully embracing this situation. But Kant's point is that this is not something that would-be thieves *can* rationally do, because the very maxim on which they act presupposes the institution of property (or something like it). There is no contradiction in the mere absence of the institution of property, rather there is a contradiction when this absence is combined with the typical intentions and likely

---

[32]Hegel (1942), §135. For discussion of this objection, see Taylor (1979), pp.78-80.

consequences associated with the maxim that the would-be thief proposes. Agents, if they are practically rational, must be capable of formulating sufficiently precise maxims and subsidiary intentions to be able to act, and this is all that provides the necessary content for the contradiction to emerge.

The contradiction in conception test is only half of the CI test, for many maxims do not involve this conceptual incoherence that would prevent them serving as universal laws, but are nonetheless impermissible. Such maxims fail the contradiction in the will test. It is here that it seems most compelling to say that Kant must really appeal to some heteronomous account of what makes the willing of a maxim impossible - for example, the participation of agents in a tradition, their possession of certain desires, or some teleological account of human nature. However, again, Kant eschews such resources because he does not suppose that the impossibility in willing resides in any psychological limitations on the agent. Instead, the impossibility resides in the logic of the interaction between the normal intentions of an agent proposing a maxim, the foreseeable consequences of the action (as with the contradiction in conception test), and the preconditions of rational agency itself.

The last of these elements is helpfully spelt out by O'Neill when she identifies five "Principles of Rational Intending" that are implicit in Kant's writings. Practical rationality demands that:

(1)We will the indispensable, or necessary, means to our fundamentally intended ends (this is the "Principle of Hypothetical Imperatives", which is found explicitly in Kant),[33] and will at least some sufficient means to such ends.

---

[33]Kant (1988), IV, 417.

(2)We seek to provide such means if they are not already available.

(3) We intend all necessary, and some sufficient, components of what is fundamentally intended.

(4) The specific intentions that we adopt in acting on a maxim must be mutually consistent.

(5) The foreseeable consequences of the specific intentions must be consistent with the maxim.[34]

These Principles of Rational Intending simply spell out the minimal requirements for rational agency. It is to be stressed that they involve no appeal to instrumental considerations concerning the satisfaction of desires. Agents are assumed to have sufficient knowledge of, and sensitivity to, their situation to form maxims on which to act, along with the subsidiary intentions appropriate to implementing this underlying principle. The five Principles of Rational Intending impose a constraint of consistency on the relationship between a maxim and its subsidiary intentions and foreseeable consequences.

When first introduced, the contradiction in the will test seemed rather vague and ill-specified - How, after all, does one generate a contradiction in the *will*? Surely such a contradiction can only be generated with the addition of something heteronomous? Indeed, Kant's own treatment, in the *Grundlegung,* of the examples of neglecting one's talents and living one's own life, not hindering, but not helping others, might seem to attract just such a reading.[35] Kant's reason for denying that a person could will, as a law of nature, that one neglect one's talents is as follows:

---

[34]O'Neill (1989), p.91-2. It would be a fruitful issue for further investigation whether there are other necessary conditions of rational agency, not necessarily to be found (explicitly or implicitly) in Kant.

[35]See, for example, Wiggins (1998), pp. 69-70.

> But he cannot possibly will that this should become a universal law of nature or that it should be implanted in us by a natural instinct. For, as a rational being, he necessarily wills that all his faculties should be developed, inasmuch as they are given to him and serve all sorts of purposes.[36]

On one reading of this, Kant has to be seen as appealing to some form of teleological account of what constitutes a rational life, according to which the possession of a developed set of talents just is one of the goals of such a life. But the introduction of the Principles of Rational Intending suggests that we are not forced to accept this heteronomous reading of Kant's (admittedly, highly compressed) route to his conclusion. That a rational being "necessarily wills that all his faculties should be developed" follows from the principle that one must will all the necessary, and some sufficient, means to one's ends (whatever these might be), combined with the requirement that the maxim on which one acts must be universalizable. The universalization of the maxim that one neglect one's talents gives the universal law, "Everyone is to neglect their talents", and in such a world no one would develop their talents, and so there would be *no* necessary or sufficient means for the pursuit of maxims.[37] The contradiction in the will consists in the fact that the agent both intends the performance of actions (because a being who failed to form any maxims at all would not be practically rational), but also does not intend any necessary or sufficient means for the performance of any actions.

Kant's other example of a case of supposed inconsistency in willing is the maxim, "Let each one be as happy as heaven wills, or as he can make himself; I will not take

---

[36]Kant (1988), IV, 423.

[37]Given that we do not live in an utterly bounteous paradise.

anything from him or even envy him; but to his welfare or his assistance in time of need I have no desire to contribute".[38] This cannot be consistently willed because:

> [A]lthough it is possible that a universal law of nature according to that maxim could exist, it is nevertheless impossible to will that such a principle could hold everywhere as a law of nature. For a will which resolved this would conflict with itself, since instances can often arise in which he would need the love and sympathy of others, and in which he would have robbed himself, by such a law of nature springing from his own will, of all hope of the aid he desires.[39]

One interpretation of this treats Kant as claiming that the duty of mutual aid is based on hypothetical prudential considerations: the duty may have costs now and then, but the benefits for me in the long run outweigh these costs.[40] But surely, it can be objected, if people who try to will this maxim are very strong and rich, they might think that there was very little chance of their ending up in a situation where they would be reliant upon the aid of others. Overall, they might consider that the risks of such a situation's arising would be outweighed by the ongoing costs of having to go to the aid of the needy.[41] John Rawls suggests that, in order to avoid this problem, a "veil of ignorance" condition has to be added to the CI test, so that the outcome of the procedure is not affected by the different risks that agents face and their differing attitudes to these risks.

But this essentially heteronomous reading gets Kant significantly wrong. In a world in which the universalized counterpart of the maxim of non-beneficence applies as a law of nature, no one would help anyone who was in need. In such a world, the ends of those who were in need and unable to help themselves would not be achieved. By the first

---

[38]Kant (1988), IV, 423.

[39]Kant (1988), IV, 423.

[40]This is clearly John Rawls's view of what is going on in Kant's argument (see Rawls (1971), p.338).

[41]For the classic expression of this objection, see Sidgwick (1981), p.389.

of the Principles of Rational Intending (the Principle of Hypothetical Imperatives) one is committed to the willing of all necessary, and some sufficient, means to any end that one might have. This presumably means that I must will that I be helped if ever I were to find myself in a situation in which I lacked the ability to achieve any of my ends by my own efforts. Thus, the proposed maxim involves me in an inconsistency, because, when I will its universalized counterpart as a law, I will that I not be helped when in need, but, as a rational being governed by the Principles of Rational Intending, I must will that I be helped when in need. The key point is that this argument does not rest on any prudential or probabilistic reasoning. It is a requirement of rational agency that I be prepared to will the means to any (permissible) end that I might have (not just the ones that I now happen to have). Finite human beings cannot guarantee that they will be self-sufficient in pursuit of their ends - this is simply a fact about *human* agency. To will as a universal law a maxim of non-beneficence is to will a world in which one of the necessary means to the ends of finite, vulnerable beings, namely help, is lacking. It is, therefore, to will the frustration of some of one's potential ends, in violation of the Principle of Hypothetical Imperatives.

But what if I am never likely to need help, surely the argument does not apply in this case? Mere unlikelihood is not enough to defeat the argument, for, as a rational being, I must will the necessary means for *any* end that I might have (so long, of course, as the willing is itself coherent). So, for the argument to fail, it would have to be, not merely unlikely that I need help, but absolutely impossible that I ever do so. Now, there might be rational beings for whom this is the case - omnipotent angels or gods, perhaps - but human beings are not such. As I said earlier, Kant is quite prepared to accept that the duties that apply to a certain species of rational being will depend on their particular abilities,

instincts, weaknesses and needs. In the case of human beings, a duty of mutual aid does exist because we are not omnipotent. Even the strongest tyrants *could* find their bodyguards slaughtered by a stronger foe or a natural disaster, even the richest tycoons *could* mislay their cheque book.[42]

So Kant's argument is very different from one based on hypothetical prudential considerations. His conclusion simply stems from the logical inconsistency involved in combining the universalized form of a maxim of non-beneficence with the Principles of Rational Intending, *given the sort of rational beings that we are*. The argument can eschew considerations about probability and risk aversion because it generates its conclusions from reflection upon the requirements of finite rational agency. However, this also means that the argument, by itself, does not go far in establishing on which specific occasions we must help others. It merely establishes that there are some occasions when we must go to the help of the needy.[43]

This is really only a cursory sketch of how the CI operates if maxims are treated as underlying intentions, but I hope that it will suffice to show that it can be made to yield some determinate results. Clearly, a proper treatment would look at how it can deal with

---

[42]In fact, tyrants would fall foul of the contradiction in conception test, because, presumably, tyrants coerce people and do not get their power by entering into free consensual relationships. Coercion fails the CI test because, if it were a universal law that everyone should coerce others, then each party complies with the will of another (when they are coerced) and simultaneously does not comply with the will of another (when they are acting as coercer). Of course, if no one were both a coercer and coerced, then some people would be coercers and others coercees, and it would be obvious straightaway that a universal law of coercion could not apply. See O'Neill (1989), p.96, and Ch.6.

[43]For a slightly different account of the duty of mutual aid, see Barbara Herman's "Mutual Aid and Respect for Persons", in Herman (1993). Herman believes that, for the duty to be established, one or both of the following must be true: there are ends that an agent wants to realize more than he or she could hope to benefit from non-beneficence, but which cannot be realized unaided, and there are some ends that no agent can forego (p.52).

a much wider range of cases and would have something to say about the role of Kant's other formulations of the CI. It might also be fruitful to explore whether there are conditions on rational agency, other than the five Principles of Rational Intending that O'Neill highlights. This could be another resource for securing determinacy in the results of the CI test.[44]

3. Towards Ethical Cognitivism

We have been examining the CI as a test of reasons for action, and we seemed to have had some success in generating determinate results that tallied with our pre-theoretic moral intuitions. The key to this success was conceiving of maxims as underlying principles or intentions. This ensured that an agent could not formulate a number of potentially conflicting maxims, as would be the constant danger if maxims are understood merely as intentions. However, the success comes at a price. There seem to be two problems confronting this approach. The first is that, if maxims were understood as underlying intentions, it will never be possible for a third party to be sure about what maxim it is that an agent acts upon and, hence, whether or not it was a morally permissible maxim. Even more worryingly, it is possible for agents themselves to be ignorant of the maxims on which they act, since maxims can be opaque.[45] The second, and related, problem is that, since this approach depends upon the maxim that an agent actually has to generate determinacy in the results of the CI test, there is no constraint on the maxims that an agent

---

[44]And given that we are not restricted to purely Kantian resources in this project, we might be able to appeal to a slightly less thin conception of rationality to find these further constraints. This would a matter for much further investigation.

[45]For presentations of this problem, see, for, example, O'Neill (1989), p.85, and Herman (1993), Ch.7, section II.

actually formulates. An agent could formulate a maxim that passes the CI test only because it is bizarrely insensitive to key features of the situation. This would be a seriously damaging flaw in an objectivist approach to ethical judgment.

But our concern is to explore the capabilities of the CI as a criterion of correct practical judgment within the wider context of a response-dependent ethical objectivism. It will be my contention that, when deployed in this role, the CI need not be so vulnerable to these problems. The reason for this is the central role that agents' sensitivities play in normatively-guided activity. It is one of the insights of O'Neill's way of understanding maxims that agents must have some situational appreciation in order to formulate the maxims that they can then test for permissibility with the CI. However, those of us who are not engaged in the interpretation of Kant, but in developing a viable ethical objectivism can go further than this. These situational sensitivities are, themselves, open to rational scrutiny - one who formulates an inappropriate maxim is guilty of a rational failure, as much as one who acts on an impermissible maxim.

But how are we to determine what constitutes a proper appreciation of the ethically salient features of one's situation? One thought might be that a maxim must be based on a true understanding of the practically salient features of one's situation. We have already seen that the ethical properties of a situation are best understood as being response-dependent.[46] So, once again, we face the question of which responses would be the appropriate ones that fix the extensions of our ethical property concepts. Again, we should start with the thought that the correct understanding of a situation as possessing a set of ethical properties E is one for which it is true that there is nothing else to think but

---

[46]I shall use the terms 'property' and 'feature' more or less interchangeably to refer to the different ethically salient aspects of an agent's situation.

that the situation exhibits the set of properties E. How do we proceed from here? One thought might be that it will not be possible to codify the reasons that compel the conclusion that there is nothing else to think, but that the situation exhibits the set of properties E. This would suggest that, ultimately, we are forced to accept an account of situational appreciation like McDowell's, even if we try to marry it to a different model of practical deliberation. There is reason to suspect that this marriage would be a turbulent one. Even if the marriage itself was harmonious, we would still face the problems that I suggested afflicted McDowell's objectivist aspirations. The prospects of this approach's yielding objective and determinate answers seem bleak.

Instead, I want to suggest a possible alternative approach, which uses the CI as a test for the rational permissibility of agents' understanding of their situation.[47] The core motivation for seeing the CI as a way of understanding the "nothing else to think" requirement was that, if there is nothing else to think but p, it must be impossible to think that not-p. The CI tests whether a principle or maxim is such that there are some members of a plurality that *cannot* act upon it.[48] One thought might be that an understanding of a situation that was not true would be one that could not be held by everyone.

If the proposal is that the CI requires that one's conception of a situation be true, one might not want to argue with this claim but one could be forgiven for thinking that we had gained little extra ground. We will have gained *some* extra ground, however. The test of correct situational appreciation will not lie with the uncodifiable sensitivities of the virtuous. We have also highlighted a commitment to the requirement that one's

---

[47]In what follows, I shall use the terms 'conception', 'understanding' and 'appreciation' interchangeably. A 'belief' will tend to mean a component of an agent's wider conception of a situation.

[48]See O'Neill (1989), p.20, and the start of Section 2 above.

conception of a situation be shareable by others, and this might be thought to suggest the sort of factors that we need to consider when we reflect upon our understanding of a situation. But, if what we are subjecting to criticism is the conception itself, it seems that we might get some rather counter-intuitive results. If a conception of a situation is one that can be shared by others, if and only if it recognises all and only the relevant ethically salient features, then we do not seem to have got further than saying that the conception has to be true. On the other hand, if the requirement is that the features that it contains must only be ethically salient ones, but that it need not contain all of these, this would be without bite against an understanding of a situation that did not report any features that were not genuinely there, but still left out much that was. Sins of omission in understanding a situation will typically be quite as serious as sins of commission.

The highly speculative proposal that I want to make is that the best way forward might be to develop recognitional norms that govern the sensitivities that we deploy, rather than the actual deliverances of these sensitivities. The Principles of Rational Intending that O'Neill thinks are inherent in Kant's work might form a model for what is involved. As was the case with the Principles of Rational Intending, the selection of these norms would be answerable to the CI, since they would have to be shareable in principle. The task would be to investigate what norms of recognition would fulfil the aim of permitting only those conceptions of situations that are, in principle, capable of being shared. The Principles of Rational Intending were understood as capable of generating contradictions when combined with a universalized counterpart of an agent's underlying maxim. In a similar way, it might be hoped that the rational impossibility of understanding a situation in a way that is guided by an illegitimate recognitional maxim would be highlighted by the occurrence of a contradiction. Agents' appreciation of their situation

would be structured by certain negative regulative recognitional principles (what one might regard as recognitional virtues).

On this proposal, the key question about the permissibility of an agent's appreciation of a situation would be whether it involved some kind of contradiction - whether there was any contradiction between the negative regulatory norms of situational appreciation holding as universal laws and everyone guiding their appreciation according to these principles. This would be a natural way of developing an account of objectivity-sustaining criticism along Kantian lines that could cope with our situational appreciation. It would also remove the threat of destabilising the project by introducing an alien account of proper situational appreciation. Of course, this proposal is enormously speculative, and it would be a considerable project to work out what these regulative norms might be and what exactly the envisaged contradiction or impossibility might amount to.[49] Ultimately, the viability of this Kantian cognitivism may depend on the successful completion of this project. It may also turn out that we have to move much further away from a strictly Kantian approach, and content ourselves with a more piecemeal attitude to the form of vindicatory explanation that will be available for judgments of situational appreciation.

If these regulative recognitional principles can do the job of delivering an objectivity-sustaining critique of our situational sensitivities, then we have taken a considerable step in the direction of cognitivism. But what model of practical deliberation do we end up with? It will, I think, be something like the following. A practically rational agent has to be sensitive to the salient features of a situation. A conception of the situation will be objectively permissible or warranted if it is the product of sensitivities that meet

---

[49]The connection with the "nothing else to think" schema should be apparent.

negative norms of rational situational appreciation. These norms will serve the purpose of filtering out ways of understanding a situation that cannot be shared by others. As with the CI test of maxims for action, the regulative norms are negative and rule out those ways of understanding a situation that are rationally impermissible. Thus, the norms do not lay down positive guidance as to what a conception must include.[50] We therefore avoid any suggestion that this model involves an aspiration for a hopelessly strong form of codification. Within these negative constraints, there may be scope for a number of different ways of understanding a situation. The possibilities of appreciating features of a situation will be potentially open-ended, and the process of refinement, guided by the regulative norms, might be an ongoing one.

In the light of their appreciation of their situation, agents formulate maxims of action. This process of deliberation is answerable for its correctness to the CI test, but the deliberation itself may not actually take the form of an application of the CI test by the agent, although it can do. Deliberation is nonetheless essential, because of the degree of under-determination that is involved in the appreciation of ethically salient features. Although it is essential that an agent be sensitive to these features, they do not, in themselves, generate reasons that have to be weighed. But an appreciation of salient features will involve an understanding of their significance for the deliberative process. An agent who does not recognise the deliberative significance of such features will not really have properly grasped their salience. In contrast to a straight cognitivist model of practical deliberation, appreciation of salient features under-determines the deliberative decision.

---

[50]Although, of course, there will be a certain minimal content that they specify.

This is, first, because these features are not themselves reason-for-action-generating and, second, because of the negative nature of the regulative norms of situational appreciation.

The focus on maxims as underlying intentions, coupled with the claim that our situational appreciation is also open to objectivity-sustaining rational criticism, aims to show how we can get determinate answers out of the CI procedure. The application of the CI test to both aspects of practical rationality offers an understanding of how the "nothing else to think" proposal could be spelt out. This is an account of correct judgment, an account of the normative theory of practical rationality, that lies at the heart of the response-dependent objectivist model. Obviously, in reality, virtuous agents will not always go through this rigamarole, they will often act instinctively, without explicit deliberation. This does not make their actions any less ethically correct, provided that the maxims that they act on would have passed the CI test. A clearer picture of this account of practical deliberation should emerge during Section 5 below, as I discuss how we might use it to address some knotty issues in the theory of practical reason.

We still have hanging over us the worry that the proposed interpretation of the CI renders opaque the moral status of an agent's maxim, and it might be thought that the explicit recognition of the role of recognitional sensitivities does not address this problem. In part, this is unfair, as it should now be clear that the appreciation of practically salient features is itself open to objective assessment. However, it might nonetheless be claimed that this assessment itself could be opaque to third parties, or even to agents themselves. Here, the fact that we have embraced a broadly Wittgensteinian account of correct concept-application should come to our aid.[51] On my interpretation of Wittgenstein, he

---

[51]See Chapters 3.6 and 3.7 above.

does not offer a constructive account of normativity in terms of communal agreement. However, the ability to use concepts in a meaningful way will nonetheless have developed for human beings (as they actually are) in a context of interaction with others and with the objects around them.[52] The ability to recognise the public conditions that warrant the use of a concept is a central part of being a competent rational being.[53] It is certainly true of a conceptually sophisticated area of our rational lives, and one that is so centrally concerned with our interactions with others, as ethics is, that the ability to deploy ethical concepts will be highly dependent on the ability to recognise publicly available features of the world. If this is indeed the case, then it makes little sense to think that the reasons that others act upon are essentially private. If they were private, then we would have few grounds for thinking that there could be rational convergence in judgment, and hence there would be no objectivity. So, from the objectivist point of view, if our maxims, whether recognitional or deliberative, were to be opaque, this would be the least of our many problems. The conditions that rendered maxims opaque would, in any case, have deprived us of the possibility for convergence in ethical judgment. On the Wittgensteinian approach, the possibility of objectivity goes along with the impossibility of opacity.

---

[52]It might be conceptually possible for a being to acquire concepts in some other way, but these would not be the sorts of concepts that we human beings have.

[53]Compare Wittgenstein's views on sensation language. That there is a central public aspect to our use of sensation language is a result of the context in which we acquire the ability to use such concepts meaningfully. The fact that sensation language involves a public manifestation does not mean that there is no such thing as pain, or that we are zombies, but that we would be unable to have developed the practice of using sensation concepts meaningfully if there were not this public aspect. See McDowell (1989) and (1991) for a useful discussion of Wittgenstein's subtle position on these matters.

4. Delivering Ethical Objectivity

We can now consider how the proposed account of ethical deliberation faces up to the challenges to ethical objectivity that we outlined in Chapter 2. In many ways, what we will have achieved will not be a definite answer, but a solution to a few problems and some clarification of the remaining issues.

In Chapter 2, I identified three very broad types of challenge to the objectivity of ethical judgments. The first challenge is that there is a fundamental incompatibility between the essentially prescriptive nature of ethical judgments and their being objective judgments about the world. Judgments are only objective if they represent a motivationally inert, neutrally-described world. I have argued that this challenge rests on a fundamental misunderstanding of objectivity. To conceive of objectivity in this way is to hanker after an illusion, which, when properly thought through, suggests an incoherent irrealism about the normative constraints on the application of concepts. On a broadly Wittgensteinian understanding of how we deploy concepts in a norm-governed manner, sensitivities and perceptions of salience play a crucial role. On this understanding, it is also hard to maintain a distinction between purely passive beliefs that reflect the world and desires that are directed at changing it. To have a belief or make a judgment can essentially involve being motivated to act in certain ways, in certain circumstances.

So the suggestion is that, if 'queerness' is a problem for ethical judgments, then it will be a problem for the normativity of any kind of judgment.[54] Even if there is a residual fear that ethical objectivism posits the existence of intrinsically motivating features of the world, it should be clear that this is not a problem from which my model of

---

[54]It is assumed that this kind of global irrealism is not something that an *ethical* anti-objectivist can willingly endorse.

objective ethical judgment suffers. The role of ethically, or more generally practically, salient properties is crucial, but these properties do not themselves generate full-blown reasons to act. A sensitive appreciation of these features merely involves a recognition that ethical deliberation should take into account certain considerations. To be sure, the recognition of ethically salient features involves a greater role for our sensitivities, interests and concerns than do judgments of some other kinds. These judgments require a more heavily conceptualized understanding of the world than do straightforwardly perceptual judgments about, for example, the shapes and colours around us. But conceptual complexity in a subject matter need not threaten its objectivity. Similarly, appreciation of the ethically salient features of the world will be more closely tied than are judgments about primary qualities to the interests and needs of human beings, and of other creatures with similar capacities for rational action and similar vulnerabilities and needs. But any problem for the objectivity of these judgments concerns issues other than 'queerness'. Reasons to act attach to the maxims that agents formulate, and these maxims are susceptible to assessment by the CI procedure for their objective permissibility. An agent's appreciation of a situation is also open to assessment for its objective permissibility, but ethical features do not generate an objective 'to-be-doneness', or anything metaphysically dubious of this sort.

So, I hope that this suggests some reasons for thinking that there need not be anything essentially problematic about the motivational role of ethical judgment. Axiological and practically rational ethical judgments both give rise to reasons, and reasons that can motivate, but they do not give rise to the same sorts of reasons. The appreciation of ethically salient properties gives one reasons to take into account certain considerations in one's deliberation. And the fully rational, and ethically sensitive, person

will be successfully motivated to do this. On the other hand, the outcome of practical deliberation also generates reasons, and it will be the mark of fully practically rational agents that they are motivated to act in the way that they have reason to do.

Let us take the other two objections to ethical objectivism in reverse order. The challenge posed by the supposed essential contestability of ethical concepts was argued to pose no independent a priori threat to ethical objectivity, unless a pressing case had been made that ethical judgments cannot command rational convergence.[55] There might be a challenge to the truth of ethical objectivism if, in fact, it turned out that ethical concepts were contested to such an extent that there was no evidence of such convergence. But such a scenario would not affect the cogency of the proposed model. It could be that this model suggests an understanding of ethical judgment that is successfully objectivist; it is just that, sadly, human beings do not deliberate, or are not capable of deliberating, in this way. This would be a little like a situation in which there were objectively correct methods of reasoning mathematically, but humans have such limited memories that they can successfully complete a calculation only very rarely. As I said in Chapter 2.6, these considerations are entirely a posteriori and are not really the job of the philosopher to assess. Anyone using these considerations to attack ethical objectivism has to offer a principled measure for the failure to reach rational convergence, an account of when such failure is genuinely objectivity-threatening, and an account of what level of it would vitiate objectivism. Even then, the most that could be shown was that, at this moment in time, and in this culture, ethical judgment cannot be regarded as being aimed at objective correctness.

---

[55]See Chapter 2.6 above.

So, the key issue that remains, and the one on which I have been most focused, is the possibility of ethical judgments' exhibiting rational judgmental convergence. In Chapter 4.4 and 4.5, I argued that objectivity, as opposed to reality, does not require the satisfaction of a width of cosmological role requirement. Nor, indeed, does it demand satisfaction of any other requirement of width of role, since width of role is, at most, a reassuring indication of reality (for cosmological role) or objectivity (for role, more broadly understood). I suggested that Wiggins's requirement that, for a judgment to be objective, there must be nothing else to think, offered a basis for a fruitful test of objectivity, provided that the requirement was construed in absolute terms. So the main question for an ethical objectivism, understood of in these terms, is whether it can offer an account of practical rationality that supports rational criticism which is sufficiently robust. The Kantian CI test seems a natural way of pursuing this strategy. In Section 2, I looked at how the CI test, understood as applying to agents' maxims for action, seemed to deliver judgments that accorded with some of our pre-theoretic moral intuitions. This gives us some confidence that it delivers recognisably *moral* judgments. The key question, then, is whether it can yield suitably determinate judgments, without bringing in ingredients from outside its theory of practical rationality, particularly elements that cannot be objectively warranted.

Understood as a test of maxims, the CI can go a considerable way in delivering determinate judgments. However, as we saw in the previous section, this development faced the charge that a focus on maxims renders it utterly opaque whether an agent acts on a maxim that is appropriate to the situation. The cognitivist move that I took sought to address this problem, and it also offered the prospect of increased determinacy. Ultimately, the question of whether the proposed account of practical rationality offers

objectivity in judgment will depend upon how successfully the project of explaining the reflective criticism of agents' situational appreciation can be executed. In particular, it will depend upon the plausibility of my suggestion that negative regulative principles of situational appreciation might be unearthed. This will be a considerable project in itself, but I hope that we have, at least, highlighted where the remaining work would need to be done to convert this approach into a truly successful model of ethical objectivism.

However, an objection might be raised that, even if this task were to be successfully completed, it would still face something like the Which Ethics? Objection that I introduced against McDowell in Chapter 5.4. The objection would be that the introduction of situational appreciation and the importance of sensitivity to salience as the basis of any rational practice mean that the reasons why it is impossible rationally to will in violation of the CI will not be absolute, but merely relative to a particular set of practices. The accusation will be that we have not escaped the threat that the supposed objectivity of ethical judgment is really only the cognoscenti marching along in step with each other. This would seem to leave us with an unpleasant dilemma: rest objective judgment solely on the CI as a test for maxims of action, and thus render the moral status of agents' maxims potentially opaque and vulnerable to the lingering suspicion of objectivity-threatening indeterminacy; or embrace the role of agents' situational appreciation and offer no genuine advance on the sort of moral realism offered by McDowell's virtue-centred approach.

The key to seeing why we can escape from this dilemma is that we accept only a weak form of uncodifiability. McDowell's approach was unable to deal with this objection because the strong Uncodifiability Thesis left it with such limited resources with which to account for objectivity-sustaining rational criticism. Unlike McDowell's theory, my

Kantian cognitivist alternative does not rely for its criterion of correct judgment on the uncodifiable situational appreciation of ideal virtuous agents. The criterion for correct judgment about reasons for action is the CI test. Although I went on to bolster this by pointing to the crucial role of agents' sensitivities in yielding determinacy for the results of the CI test, these sensitivities are understood in a radically different way from that in which McDowell understood them. The criterion of the correct appreciation of the ethically salient features of one's situation is not uncodifiable virtues, but one's appreciation being objectively permissible - its being a conception of the situation that anyone can hold, a conception derived from sensitivities that are governed by regulative principles of rational situational appreciation. This does not import any illicit codification because, as with maxims of action, these norms are purely negative. Crucially, although the importance of situational appreciation is recognised, the criterion for correct situational appreciation is not definitionally clamped to the appreciation of any group, ideal or otherwise.

One way of understanding McDowell's objectivist ambitions is to see him as embracing Donald Davidson's idea that relativism makes no sense because we can find "no intelligible basis on which it can be said that [conceptual] schemes are different".[56] The implicit aspiration is that his position avoids relativism because it literally makes no sense to think of there being other rational ethical schemes. Judgments other than those of the virtuous would not be (fully) interpretable as moral, and hence would not genuinely be in conflict with those of the virtuous. Thus, the failure of the virtuous to be able to

---

[56]"On the Very Idea of a Conceptual Scheme", in Davidson (1984), p.198. For more on Davidson's rejection of the possibility of different conceptual schemes, see also Davidson (1985a), (1985a), (1989a) and (1989b), and Child (1993-4).

offer comprehensible justifications to others does not amount to dogmatism. As I mentioned in Chapter 5.4, this sort of idea might have some plausibility if we regarded the virtuous form of life as wide and non-exclusive. But then the Uncodifiability Thesis makes an appeal to the uncodifiable responses of the virtuous (conceived broadly) utterly vacuous as a criterion of correct ethical judgment. If, instead, the virtuous are regarded as an elite group, the mooted Davidsonian move seems wildly dogmatic and implausible.

In contrast, our Kantian cognitivism can make the Davidsonian move in order to defend itself from the charge of relativism, without being guilty of dogmatism. This is because it is not simply the fact that a judgment is part of the practice of the virtuous that makes it morally correct, but that its objective permissibility has been licensed by the CI test. Only those who really lack any ethical, or practically rational, concepts would be unintelligible, and thus fall outside the practice of practical reason. But these beings cannot be regarded as possessing another, rival viable ethical scheme, they simply have no ethical scheme at all.[57] Thus, we seem able to achieve objectivism without relativism or dogmatism, in a way that McDowell could not.[58]

What I have said so far does not demonstrate that ethical objectivism is true. It does, I hope, offer a plausible sketch of how ethical judgment might take a form that is open to objective assessment, and, perhaps more importantly, it highlights what more needs to be done to advance the discussion. However, we can, I think, say at this stage that, if the details of the model are successfully filled out, the resulting approach will not

---

[57]See, also, the discussion of the categorical nature of moral requirements in Section 5(b) below.

[58]For a useful discussion of how the Davidsonian move can avoid the charge of conceptual chauvinism or dogmatism, see Hurley (1992). Jonathan Lear, in Lear (1982) and (1984), interprets Wittgenstein as offering an idealistic view of reality that would seem to be susceptible to the charge of dogmatism (see Stroud (1984) and Seabright (1987) for challenges to Lear).

just be objectivist, but also cognitivist. The cognitivism, as opposed to objectivism, will be of a slightly limited form. The CI is used as a tool for the reflective criticism of our appreciation of the ethically salient features of our world, and of the recognitional sensitivities that we require to achieve this appreciation. It can thus be claimed that, if one comes to a correct appreciation of one's situation, then one has come to believe something for which it is appropriate to say that there is nothing else to believe. The criticism supplied by the CI thus ensures that, where there is convergence in judgment and the judgment passes the test, the convergence can be truth-sustaining. Since we can now regard some of our valuational beliefs as being true or false, we have a form of ethical cognitivism.

The form of cognitivism will be fairly weak for two reasons. The first is that the CI is a test of objective permissibility, so it may be that there will be different permissible ways of understanding situations, or subtly different permissible maxims of sensitivity that give different understandings of a situation. So, the features of a situation that any permissible appreciation must reflect may not cover every detail. But this under-determination seems to be handled in a rather healthy way. The CI is a recursive test of rational permissibility, and the process of rational reflection will be an ongoing one, offering the possibility that this under-determination need not be a brute feature of our ethical world that simply has to be accepted. This form of cognitivism also eschews the fantastic ideal of a perfectly sensitive agent. There are no perfectly sensitive ethical beings (nor, indeed, are there beings who are perfectly sensitive to any other kind of property), and a model of ethical appreciation that does not labour under this false hope is surely to be welcomed. It seems to offer a form of ever-modest optimism in the face of the under-determination of our ethical appreciation.

The second reason that this cognitivism will be of a weak variety is the relationship between ethical appreciation and practical deliberation. Ethical appreciation is not ethical judgment, but offers the basis on which that judgment can be made. Correct practical judgment takes the form of deliberation in formulating a maxim, not the acquisition of beliefs about one's situation, nor of beliefs about what is required. When the CI assesses maxims of action, it issues in an assessment of the objective permissibility of what one is to will, not of what one is to believe. Thus, deliberation about how one should act is essentially practical. Of course, we can naturally say that, when we have assessed a maxim, we come to know something, namely, for example, that it is objectively permissible to ϕ in situation S. One acquires a reflective belief about one's reasons. If these reasons pass the CI test, then they command rational convergence. Therefore, we can, quite innocuously, regard the beliefs that we have about our objectively warranted reasons as being true. We can thus say that we know that it is permissible to ϕ in situation S. But we should not forget that this cognitivism is a secondary phenomenon. The real work in deliberation is done by critical reflection that is objectivist in nature, but not cognitivist. In the case of the output of practical deliberation, the cognitivism is doubly weak since it is also affected by the under-determination that I suggested applied to valuational appreciation.

5. The Payoff

There still remains much more to be said before we have a satisfactory understanding of the interaction of the CI with our ethical sensitivities, which yields both objective correctness in our appreciation of the ethically salient features of situations and truly determinate reasons for action. Only when this process is more complete will it be possible

to say whether the proposal offers a genuinely plausible account of how agents' valuational and action-guiding judgments can be objectively correct. However, we have, I hope, reached a point where we have a reasonably clear sketch of what correct practical deliberation might look like on this model. So, before drawing this chapter to a close, I want to go one step further in spelling out the advantages that a Kantian cognitivist account of correct ethical judgment might offer. In this section, I shall try to sketch how this model might be developed to offer some fruitful ways of addressing some tricky and controversial issues in moral psychology and meta-ethics. The issues are: the particularity of ethical judgment; the categorical nature of moral reasons; and weakness of will, or akrasia. It should be stressed that these suggestions really are only advertisements for how the project might be developed - they are even sketchier than what has gone before. The success with which the account deals with these issues will, ultimately, be entirely dependent upon how the details of the central part of the model are filled out. However, these prospective benefits might give further reasons for thinking that it would be worthwhile to investigate the approach more extensively. Exploring these benefits also sheds more light on how the model might operate.

(a) The particularity of ethical judgment

One of the most common accusations levelled at Kantian ethics is that it ignores the particularity of our moral judgments. With its supposed emphasis on rules of reason and rigid, categorical duties, Kantian ethics ignores the particular complexities of the situations of moral decision that we face, and tries to cram these complexities into a rigid

framework.[59] It is often claimed to be one of the great advantages of the virtue-centred approach to ethics that its emphasis on the situational appreciation of the virtuous and the uncodifiable nature of practical rationality means that it can fully respect the particularity of each situation.

It should already be clear from my discussion that, for a number of reasons, this objection will be harder to sustain against the Kantian cognitivist position that I have suggested. The first point to make is that the role of the CI is not to lay down a set of rigid duties. Its role is as a mechanism of negative recursive criticism of the maxims of action and situational appreciation of ethical agents. It gives no positive guidance, but rules out what it is rationally impermissible to do or believe. Far from being too rigid in its prescriptions, the more pressing danger seems to be that it might leave too open what count as permissible courses of action for it to deliver genuine objectivity.

However, it might be thought that, even though the CI is a negative regulative principle, it is nonetheless too crude in the way that it rules out courses of action as impermissible, since its universalism means that it must ignore genuine ethical features of situations. But, again, this accusation seems unsuccessful because of the role that ethical features, and our sensitivity to them, play in practical deliberation. Because of the approach's Wittgensteinian background, the rational deployment of concepts is seen to depend on the sensitivities and perceptions of interest and salience of deliberating subjects. In order to use ethical concepts properly, agents must have the required sensitivities to be able to form a proper understanding of a situation. Because the CI does not generate duties out of nothing, but tests agents' maxims for action, agents have to be suitably

---

[59]See, for example, Winch (1965), and the discussion of this article in Wiggins (1998), pp.166-82.

rationally competent and sensitive to their situation to be able to formulate maxims of action. So, the first point that we can make against the particularist challenge is that this approach allows the agent whatever sensitivity to a situation is required to form a maxim of action. The deployment of the CI takes an agent's maxim as it is, and the required degree of sensitivity of the maxim to the particularity of the situation will vary from case to case.

However, a different problem might now seem to emerge. It might be conceded that the model allows whatever sensitivity to the particularities of a situation that is possessed by an agent's maxim to be transferred to a judgment about what is morally permissible. But it might then be objected that this still allows genuine *ethically* salient features of a situation to be ignored, because they might be ignored by the agent when the maxim is initially formulated. The problem for a Kantian approach that avoids rigorism by taking an agent's maxim to be as particular as necessary is that it seems then to have no resources for commenting on the appropriateness of a maxim to its situation. At this point, we must not forget the cognitivist aspect of our proposed model of practical deliberation. It is not enough that agents form maxims of action that are sufficiently particular to the features of their situation for their own purposes. Agents also have to be sensitive to the features of the situation, and their understanding of the situation has to be able to withstand reflective criticism. A conception of a decision situation that ignores ethically salient features will not be one that every rational agent thinking about the situation could accept, in other words the conception will fail a CI test for situational appreciation. It will not, therefore, be objectively permissible. So, because of the double aspect nature of the model, a correct ethical judgment will be sensitive to any ethically salient features - the

proposed maxim of action will both pass the CI test and be based on an ethical conception of the situation that is objectively permissible.

Of course, this is, to an extent, an idealization. What we have specified here is ethically correct judgment, and we will often fail to achieve this because we do not have a perfectly correct appreciation of the salient features of our situation. But this idealization is innocuous for two reasons. The first is that the challenge was that a Kantian approach gives an account of (ideally) correct deliberation that ignores salient features of situations. I have, I hope, suggested that the situation is entirely the reverse - it is a requirement of practical rationality that one be perceptive to all ethically salient features. The second reason is that the idealization is not extreme. This is because of the dual aspect nature of our account of practical deliberation. Agents may formulate maxims that pass the CI test (that is, they are maxims of action that anyone can act upon), even if their conception of the situation does not strictly pass the test. There may be ethically salient features of the situation that the conception ignores, but the maxim is nonetheless one that could also be formulated by one who had a permissible appreciation of the situation. I rather suspect that this describes our usual decision-making. Because the CI test is a *negative* test of permissibility there may be a number of different ways of understanding a situation that are permissible, and there may be a number of different maxims that are permissible reasons for action based on a particular situational appreciation. It is possible that an agent can, in practice, have the latter without having the former.

To push the objection about particularity any further than this would seem to presuppose the strong Uncodifiability Thesis, according to which Wittgenstein's rule-following considerations establish the conceptual primacy of a subject's situational appreciation over norms or rules. But the Weak Uncodifiability Thesis, which seems to

have a genuine Wittgensteinian provenance, is quite compatible with our understanding of rational transactions within our cognitive practices making essential appeal to metaphysically modest rules or norms.

(b) The categorical nature of moral reasons

One of our pre-theoretic intuitions about morality is that moral demands have a categorical status, that is, they are not dependent for their reason-giving force on any extra-moral element, such as an agent's possession of a certain desire or inclination. How might the proposed approach deal with this issue?

In order to approach this issue, I want to consider what McDowell's model of practical rationality can say about categorical moral reasons, and then go on to see how we might have the resources to improve on this. I mentioned, in the previous chapter, John McDowell's suggestion that the explanation of virtuous action need not involve reference to an independently intelligible desire or interest.[60] Although I raised some serious objections to the cogency of McDowell's model of deliberation, we can still make use of this proposal. On my model, the ideal of ethical deliberation is an agent formulating a maxim of action that passes the CI test and is also based on an appreciation of the situation that is objectively permissible. That the maxim to ϕ passes the test and is based on a permissible appreciation of the situation means that it provides a normative reason for the agent to ϕ. This simply follows from our account of practical rationality. In Chapter 4.2 and in Section 4 of this chapter, we saw how, on the Wittgensteinian conception of objectivity, the correct application of a concept may not simply require one

---

[60]See Chapter 5.5.

to believe or infer certain things, but may also require one to act in certain ways. Thus, to be the sort of person who correctly judges that p one might also have to be the sort of person who is motivated to act in certain ways, when the circumstances demand. From this we can claim that, to the extent that agents are rational, they will be motivated to ϕ if they have a normative reason to ϕ.

This describes a successful case of correct ethical deliberation and action based on it, and it uses materials supplied only by our conception of practical rationality - there is no appeal to independently intelligible desires or other extra-moral elements. We can also adopt McDowell's disjunctive account in order to explain cases of immoral action, although things will be a little more complex than they were on McDowell's approach. There are a number of things that can go wrong in the process of ethical deliberation. An agent can have a faulty appreciation of a situation, and thus will an inappropriate maxim based on this conception. Or an agent can have a permissible appreciation of the salient features of a situation, but will a maxim based on it that fails the CI test. Or an agent can appreciate the situation aright and will a permissible maxim, but fail to be effectively motivated to act. The explanation of what goes wrong in ethical deliberation may make essential reference to desires (for example, because an agent's desires figure too prominently in his or her appreciation of the situation), but it need not.

For Kant, actions that are immoral are also irrational. Many have seen this as amounting to the claim that a charge of irrationality can appropriately be levelled at those who are deaf to categorical moral reason, and they regard this claim as a mere 'bluff'.[61] McDowell accepts that it would amount to a 'bluff' to appeal to the categorical nature of

---

[61]See, for example, "Internal and External Reasons", in Williams (1981), and Foot (1972). The term 'bluff' is Williams's (Williams (1981), p.111).

moral requirements to condemn the immoral as irrational; however, he nonetheless regards moral requirements as categorical. In McDowell's view, it does not amount to a 'bluff' because it is necessary to appeal to some desire or interest that agents possess in order to give someone a reason to act. Rather, it amounts to a bluff because one who lies outside the virtuous form of life will be insensitive to the relevant features of situations, and thus insensitive to the charge of irrationality.[62] Someone who stands outside reason-giving practices cannot be reasoned with, but has to undergo some sort of non-rational induction into these practices (of the sort that a child undergoes in the early stages of education, or an adult undergoes during a religious conversion experience). On this view, the most that can be said of such people is that, purely as a matter of description, they are acting non-rationally, and this must not be seen as a claim that is likely to sway them to act rationally.

But McDowell seems to concede too much to the critic of categorical reasons. Clearly, it makes no sense to try to offer reasons to someone who genuinely lies outside *all* rational, reason-giving practices, but this does not mean that we need to abandon the claim that one who acts immorally acts irrationally. McDowell seems to be drawn to make this concession because of his theoretical reliance on the uncodifiable sensitivities of virtuous agents. For those who are close to achieving this ideal, a failure of ethical appreciation on their part might be corrected by suggestions, hints and useful analogies, but for anyone beyond this charmed circle nothing short of induction into the practice will do the job. For those who are within the charmed circle of the virtuous, ethically correct action requires no other explanation than the agent's possession of an accurate conception of the situation.

---

[62]See McDowell (1995b), p.75.

On my model of practical deliberation, we can go rather further than this. We can accept that, for those who really do lie outside all ethical, or even all rational, practices, reasoning with them will not help them to improve their ethical judgment (because such people do not make such judgments in the first place). However, such cases are likely to be extremely rare. On the Kantian cognitivist account, there are various points at which agents can suffer a rational failure to deliberate properly - they can have a faulty appreciation of the practically salient features of their situation, they can will a maxim that fails the CI test, or they can will an appropriate maxim, but fail to be properly motivated to act. The way in which correct judgment features both at the level of sensitivity to the practically salient features of a situation, and at the level of action-oriented deliberation, means that there need be no temptation to regard those who deliberate in a faulty manner as beyond the rational pale. Because correct judgment is not understood in terms of uncodifiable sensitivities, but in terms of permissibility according to the CI test, it is possible to cite reasons why an agent failed to deliberate properly. The failure to achieve an accurate appreciation of one's situation is not a failure to be sensitive in the uncodifiable manner of some idealized virtuous agent, but is a failure to recognize features that would be included in an objectively permissible conception of the situation. There is no reason why agents cannot have pointed out to them features that they have ignored or overemphasised. And the very form of a failure to be guided by maxims of rational situational appreciation will also suggest the way in which an agent might be reasoned into improving their appreciation of the situation. Only in extreme cases would such failure be a result of a complete lack of sensitivity.

In the case of deliberation about a proposed maxim, it is even more clear that a charge of irrationality might be in order when this deliberation goes wrong. Because the

criterion of correct deliberation is permissibility by the critical reflective test of the CI, a mistake in reasoning can be pointed to when an agent deliberates incorrectly. Since the permissibility of maxims is under-determined by one's appreciation of a situation, there is much less scope for a failure of deliberation to be the result of a brute lack of sensitivity. So, I hope, this suggests a way in which the Kantian cognitivist model might start to provide a plausible way of understanding the claim that a failure to act on a categorical reason can constitute a genuine failure of rationality.

(c) Akrasia

David Wiggins writes of akrasia, or weakness of will, thus:

> Almost anyone not under the influence of theory will say that, when a person is weak-willed, he intentionally chooses that which he knows or believes to be the worse course of action when he could choose the better course; and that, in acting in this way the weak-willed man acts not for no reason at all - that would be strange and atypical - but irrationally.[63]

Cognitivist accounts of ethical deliberation have a particular problem with this phenomenon. For there is a simple argument that, if one who acts virtuously and one who acts akratically both recognise that they have a reason to ϕ, but the akratic fails to ϕ, then there must be some motivational state that the former possesses and the latter lacks. Since the two agents' reasons, and thus cognitive conceptions of the situation, are supposed to be the same, this motivational element must be purely non-cognitive.

As we saw in Chapter 5.5, McDowell's response was to deny the premise of the argument - the akratic and the virtuous do not share a common conception of the situation; the akratic shares with the continent person a conception of the situation that

---

[63]"Weakness of Will, Commensurability, and the Objects of Deliberation and Desire", in Wiggins (1998), p. 239.

is clouded by desire, but lacks the continence to overcome this defect. We saw, in Chapter 5.6, that there are reasons to be suspicious of McDowell's claim that idealized virtuous judgment involves the 'silencing' of counter-moral considerations. This approach also has the weakness of suggesting that, for akratics, the problem is that they are not properly self-aware - that is, they fail to recognise that they have a clouded perception of the situation. But in fact the fault that akratics suffer from is not that they misunderstand themselves, but that they *do* what they judge to be inferior. McDowell's treatment of akrasia would also seem to be afflicted by the flaw that Wiggins identifies, namely being "under the influence of theory".

So, the question that I shall very briefly address is whether our Kantian cognitivist model of deliberation might offer a better response to the challenge posed by the phenomenon of akrasia. There seem to be reasons to suspect that it might. The key to this is that the model envisages several different aspects of ethical deliberation, and akrasia (and other irrationalities) could occur when these come out of alignment. One way of explaining the phenomenon might be to focus on the link between agents' appreciation of normative reasons and their motivation. I said that it might be possible for a person to form an accurate conception of a situation and formulate a maxim that passes the CI test, but then fail to be motivated to act in this way. As it was claimed that it can be part of correctly applying a concept that one be motivated to act in certain ways, agents who fail to ϕ when they have a normative reason to do so would count as acting irrationally. This might be thought to explain how agents can genuinely judge that they have most reason to ϕ and act irrationally when they not-ϕ. However, this suggestion does not really deal with the phenomenon. There clearly is some failure of rationality on the part of agents who fail to act on what they have most reason to do, but this proposal does not explain

how, when they ψ, rather than ϕ, they nonetheless act for a reason.[64] The fault here does not seem to be akrasia but lack of the executive virtue of continence (which may not be quite the same thing).

However, we might have available to us a much better explanation of akrasia. The key to this is the fact that the conception of a situation that an agent develops can (and usually does) under-determine what course of action is practically correct. This under-determination occurs because, first, the conception is tested only to see if it is objectively permissible (thus introducing one level of under-determination), and, second, there may be more than one permissible maxim that an agent can formulate on the basis of this conception. Also, of course, the features of a situation, though practically relevant, are not themselves reason-for-action-generating. It is possible for an agent to formulate an appropriate maxim, that is, come to a correct judgment about what action is permissible, but nonetheless have an inadequate conception of the practically salient features of the situation. But the agent's conception of the situation can support other, morally impermissible maxims for action. The akratic might be regarded as someone who has an inadequate conception of the situation, nonetheless formulates a permissible maxim, but then is motivated to act according to an alternative, impermissible maxim. In fact, the akratic does not necessarily have to have a faulty conception of the situation. The under-determination of deliberation by situational appreciation means that an agent can act for a reason, in accordance with one of the impermissible maxims based on a perfectly adequate conception of the situation.

---

[64]The phenomenon that this proposal explains might seem to be closer to accidie than akrasia.

This sketch of an explanation of akrasia seems to have some attractive features: akratics really do judge that ϕ-ing is the right thing to do; they act irrationally because they fail to be motivated by a normative reason that they recognise; and, yet, they also act for a reason, because they act on a maxim based on their faulty appreciation of their situation. Also, we seem to be able to account both for those cases of akrasia where there is some plausibility in following the Aristotelian line that the akratic's conception of a situation is faulty or shallow compared with that of the truly virtuous, and those cases where the akratic really does seem to reach the same practical judgment as the non-akratic, or even the virtuous.

Although we have discussed akrasia as a failure to do what is judged to be morally required, the model of practical reason, according to which features of situations are not themselves reason-for-action-generating and deliberation is separate from, and under-determined by, the appreciation of these features, could be disentangled from the CI. Thus, one could use this as the basis for a model of akrasia in non-moral cases, that is a failure to do what one recognises to be practically rational, where practical rationality is not equated with morality.

So, the model of deliberation that I have suggested might allow us to rebut the non-cognitivist's attempt to use the phenomenon of akrasia as an argument against ethical cognitivism. It might also allow us to offer a rather less counter-intuitive understanding of akrasia than McDowell can offer. Because deliberation is separate from situational appreciation, we are not forced to claim that ideal deliberation involves a blindness to certain situational features. Since situational appreciation radically under-determines deliberation, and since practically relevant features are not themselves reason-for-action-

generating, we can accept the pre-theoretic idea that the akratic really does judge correctly what is required.

6. Some Concluding Remarks

The time has come to try to draw together the strands of my argument. My aim has been to set out and defend a variety of ethical objectivism. It has been my contention that the 'Kantian' cognitivism that I have proposed provides a potentially fruitful way of treating judgments about what one ought to do as objectively correct or incorrect, and judgements about the ethical features of situations as true or false.

The thesis has taken the form of an attempt to circle in on a set of conclusions, and, in many ways, this process of circling in still has a very long way to go. But let us review the results so far. In order to give the opponents of ethical objectivity their best chance, we started out by surveying some of the broad strategies that ethical non-cognitivists or anti-objectivists have pursued. They can accept that ethical discourse has objectivist or truth-stating ambitions, but claim that it fails to fulfil these ambitions. They can claim that the surface features of moral language may seem truth-apt, but that really, deep down, moral discourse has some expressive, or other non-cognitive, function. Or they can accept that moral discourse is truth-apt, but claim that the variety of truth predicate that is in use is not genuinely objectivity-sustaining. I noted that each of these strategies has problems of its own, but that the success or failure of any one of them, in itself, would not make ethical objectivism or cognitivism any more tenable.

Instead, I attempted to draw out the key problems that all of these various strategies saw as imperiling the objectivity of ethical judgment. These problems concerned the relationship between ethical judgment and action, the lack of proper rational

convergence in ethical judgment, and the supposed essential contestability of ethical concepts. I attempted to show that the third of these problems depended on the second for any a priori force that it might have, and that the first two problems gain much of their power from a certain conception of objectivity. This conception of objectivity takes the objectivity of scientific discourse as a paradigm and suggests that judgments can only be objectively correct if they are responsive to features of the world that are entirely independent of any perspective or of the particular sensitivities of observers. According to this view of objectivity, the essential relationship between ethical judgments and the reasons that agents have to act, and the fact that ethical judgments are not responsive to features of a neutrally describable world, fatally undermine the objectivity of these judgments.

Attractive though this conception of objectivity might be, it is of dubious coherence. If the meaning of concepts is not objective - if there can be no objectively correct application of a concept - then we would face a situation of global irrealism. Ethical anti-objectivists cannot accept such a possibility, since it is essential to their position that ethics fails to match up to a standard that other domains - paradigmatically, physical science and mathematics - do successfully satisfy. In Chapter 3, I drew on arguments suggested by Wittgenstein's consideration of rule-following to argue that, if the strongly metaphysical or philosophical realist conception of objectivity, embraced by ethical anti-objectivists and non-cognitivists, is thought through, then we end up in a position of meaning irrealism. A platonic realism about meaning is found to be hopelessly incoherent and mythological, while a dispositionalist reductionism cannot, ultimately, explain the normativity of meaning. Meaning irrealism was also seen to be of very dubious coherence. The conclusion that I drew from this was that there was something seriously

wrong with attempts to provide an external validation (or, indeed, invalidation) of normative practices in terms of the non-normative. Wittgenstein's own non-reductionist, 'common sense realist' account of rule-following stresses the importance of participation in irreducibly normative customs or practices. It seems to offer a genuinely objectivist account of meaning, whilst eschewing platonism, and hyper-objectivist metaphysical or philosophical realism more generally.

My examination of Wittgenstein's rule-following considerations supplied what are, arguably, compelling reasons for questioning the coherence of what ethical anti-objectivists demand of putatively objective ethical judgments. The inextricable role of basic actions and subjects' sensitivities and interests in the proposed account of the objectivity of meaning suggests that it is much too hasty to complain that ethical judgment cannot be objective because of its links to action and motivation. The rejection of platonism and reductionist accounts of meaning suggest that the demand that objective judgments reflect features of a neutrally describable world is also in doubt. So, the first of the three aims that I set myself at the start of the thesis has been met: we have found some good reasons to reject the ethical anti-objectivist's underlying conception of reality and objectivity.

However, it would also be too quick to suggest that this means that ethical objectivism is shown to be true. There are a number of domains that are marked by apparently truth-apt language that do not exhibit much, or, indeed, any objectivity. If we have done nothing to show that ethical judgments might be any more objective than judgments about, say, the humorous or the tasty, then we have not really defended the objectivity of ethics. In Chapter 4, I suggested that the appropriate way of addressing this problem was to think of ethical concepts as response-dependent. A response-dependent

conception of a domain can support objectivism, provided that the judgments in the domain meet a suitably robust standard of rational criticism. I rejected the suggestion that judgments must meet some width of cosmological or causal role requirement if they are to be genuinely objective. Such a requirement seems to be neither a necessary nor a sufficient condition for objectivity. In any case, a rigorous interpretation of such a requirement seems to be a product of the illicit conception of objectivity that we have rejected, and a more liberal interpretation seems to allow for the objectivity of ethics. Instead, I suggested that a domain is genuinely objective to the extent that its norms of reasoning leave, in an absolute sense, no rational option but to reach a certain judgment. A conceptual practice cannot simply secure its objectivity by insulating itself from other practices. The key question for ethical objectivism is now one of whether ethics supports norms of rational reflection and criticism that are suitably demanding.

A natural starting point for addressing this question was John McDowell's virtue-centred account of practical rationality, since he also sees a key argumentative role for Wittgenstein's rule-following considerations in debates about ethical objectivity. However, on his view, the rule-following considerations do not simply undermine the non-cognitivist's knock-out argument, they also support a degree of uncodifiability that places considerable theoretical constraints on how practical deliberation is to be understood. The Uncodifiability Thesis, which McDowell seems to advocate, rules out an approach to the question, "How should I live?", other than via the concept of the virtues. In Chapter 5, I argued that this strong Uncodifiability Thesis is not, in fact, supported by Wittgenstein's arguments. The rule-following considerations can only be seen to support a Weak Uncodifiability Thesis, which excludes attempts at a metaphysically ambitious codification of rationality. I also suggested that there are conceptions of the role of rules, norms or

principles in ethics (namely, rules as constitutive of practices and negative norms of recursive criticism) that seemed to be entirely compatible with the Weak Uncodifiability Thesis and thus with the Wittgensteinian conception of objectivity that I have advanced.

This, of course, does nothing to rule out the possibility that a virtue-centred approach to response-dependent ethical objectivism, such as McDowell's, is nonetheless the right one to pursue. However, I went on to argue that McDowell's virtue-centred approach was seriously flawed as a way of understanding the theory of practical rationality that plays such a key role in a response-dependent conception of ethical objectivity. Because of the restrictions imposed by the strong Uncodifiability Thesis, it is hard to see how internal criticism of judgments could be genuinely objectivity-sustaining. McDowell's virtue-centred approach seems condemned either to dogmatism or to relativism. We also found that the strong Uncodifiability Thesis seems to ensure that the approach either has no model of proper deliberation, or else has one that is so hopelessly idealized that it can play no useful role as a normative guide to proper reasoning. Again, this is a serious lacuna in any proposed response-dependent objectivism. Thus, the second goal that I set myself has been met: we have found reason to suppose that a response-dependent, virtue-centred approach to ethical objectivity is neither required by the Wittgensteinian conception of objectivity, nor is itself adequate as an objectivity-sustaining account of practical rationality.

The fact that Wittgenstein's arguments do not support the strong Uncodifiability Thesis is fortunate if the weaknesses in McDowell's positive account are as serious as I have suggested. We now have much more conceptual space in which to hunt for alternative non-reductionist ways of spelling out the theory of practical rationality that determines the objective correctness of ethical judgments. In this last chapter, I have tried

to spell out how a Kantian account of practical rationality might function in this role. Using the CI not as a generator of rigid duties, but as a norm of recursive rational criticism, I have tried to show how it can establish the objective permissibility of agents' reasons for action. The CI is part of a non-reductionist, non-foundationalist approach to rationality and imports no illicit codifiability, since it is a purely negative recursive principle. It thus seems to fit very naturally into the broader Wittgensteinian conception of objectivity with which I have been working. The Wittgensteinian conception of objectivity, in which subjects' sensitivities and perceptions of salience play such a crucial role, also allows us to see that the CI might be used as a key element in a process of reflective rational criticism of the ethical beliefs that agents have.

This model of practical rationality might, thus, not only support an objectivism about agents' reasons for action (as Kantian ethics has traditionally attempted to do), but also offer an objectivism about agents' appreciation of value. A fruitful line of investigation would be to see what maxims of rational situational appreciation there might be that govern the permissibility of subjects' appreciation of their situation. The selection of such maxims would itself be susceptible to the CI test to ensure that the maxims were genuinely shareable and hence capable of sustaining objectivity in situational appreciation. Ethical beliefs that stand up to this kind if scrutiny would have a high enough degree of rationally warranted convergence in judgment to count as not only objectively correct, but also true. We would thus have a form of ethical cognitivism, but one with a much more plausible theory of correct practical deliberation than a virtue-centred approach that relies only on virtuous situational appreciation. The more explicitly cognitivist aspect of the model offers a way of avoiding the problem that Kantian theories may remain silent on the way that agents can formulate (and test) inappropriate maxims on the basis of an

inadequate appreciation of the ethically significance of features of their situation. Agents' appreciation of the ethically salient features of their situation is also answerable to standards of objectivity-sustaining rational criticism.

Ultimately, the process of circling in on a plausible objectivist account of practical rationality has not got very far. We have rebutted a key anti-objectivist assumption and shown that this rebuttal does not land objectivists in more trouble by foisting on them an inherently problematic conception of practical rationality. We have also sketched a proposal for what a more fruitful objectivist conception of practical rationality might look like. I hope that I have said enough to suggest that this approach has a certain theoretical coherence and potentially offers an account of practical deliberation with some attractive features. There still remains a huge task to be undertaken to show that this approach can deliver objectively warranted answers about what to do in a much wider range of cases than I have considered here. In particular, the interface between the CI test and our ethical sensitivities has barely been touched upon. A fully satisfying development of this approach would need to say much more about the regulative norms of rational agency that determine objectively warranted ethical sensitivity. Only if this area of investigation can withstand further scrutiny will the model really have succeeded in explaining how we can have valuational ethical judgments that are true, and reason-giving judgments that are sufficiently determinate to yield genuine ethical objectivism. Pending this development, we have still learnt something about the potential of response-dependent ethical objectivism and the true limits imposed by the uncodifiability of rationality - the task of formulating and defending a convincing ethical cognitivism is far, far from complete, but the possibilities for pursuing this task are much more extensive than they might have appeared at one stage in the argument.

BIBLIOGRAPHY

ARISTOTLE (1985), *Nichomachean Ethics,* Terence Irwin, translator (Indianapolis, Hackett).

AYER, A.J. (1971), *Language, Truth and Logic* (Harmondsworth, Penguin).

BAKER, G.P. and HACKER, P.M.S. (1984), "On Misunderstanding Wittgenstein: Kripke's Private Language Argument", *Synthese*, 58.

BAKER, G.P. and HACKER, P.M.S. (1989), *Wittgenstein: Rules, Grammar and Necessity* (Oxford, Blackwell).

BLACKBURN, Simon (1971), "Moral Realism", in *Morality and Moral Reasoning*, John Casey, editor (London, Methuen, 1971).

BLACKBURN, Simon (1981), "Reply: Rule-Following and Moral Realism", in Holtzman and Leich, editors, *Wittgenstein: To Follow a Rule.*

BLACKBURN, Simon (1984a), "The Individual Strikes Back", *Synthese*, 58.

BLACKBURN, Simon (1984b), *Spreading the Word* (Oxford, Clarendon Press).

BLACKBURN, Simon (1985), "Errors and the Phenomenology of Value", in Honderich, editor, *Morality and Objectivity.*

BLACKBURN, Simon (1993), *Essays in Quasi-Realism* (Oxford, Oxford University Press).

BOGHOSSIAN, Paul (1989), "The Rule-Following Considerations", *Mind,* 98.

BOGHOSSIAN, Paul (1990), "The Status of Content", *Philosophical Review*, 99.

BUDD, Malcolm (1984), "Wittgenstein on Meaning, Interpretation and Rules", *Synthese*, 58.

BUDD, Malcolm (1989), *Wittgenstein's Philosophy of Psychology* (London, Routledge & Kegan Paul).

CASSAM, Quassim (1986), "Necessity and Externality", *Mind,* 95.

CAVELL, Stanley (1969), *Must We Mean What We Say?* (New York, Charles Scribner's Sons).

CHARLES, David (1995), "Aristotle and modern realism", in Heinaman, editor, *Aristotle and Moral Realism.*

CHILD, William (1993), "Anomalism, Uncodifiablity and Psychophysical Relations", *Philosophical Review,* 102.

CHILD, William (1993-4), "On the Dualism of Scheme and Content", *Proceedings of the Aristotelian Society*, 94.

CHILD, William (1994), *Causation, Interpretation and the Mind* (Oxford, Clarendon Press).

CONANT, James (1996-7), "On Wittgenstein's Philosophy of Mathematics", *Proceedings of the Aristotelian Society*, 97.

CORNWELL, James (1995), *Objectivity, Practice and Ethical Value* (University of London, M.Phil. Thesis).

CRAIG, Edward (1975), "The Problem of Necessary Truth", in *Meaning, Reference and Reality*, Simon Blackburn, editor (Cambridge, Cambridge University Press, 1975).

CRAIG, Edward (1984), "Arithmetic and Fact", in *Exercises in Analysis*, Ian Hacking, editor (Cambridge, Cambridge University Press, 1984).

DANCY, Jonathan (1993), *Moral Reasons* (Oxford, Blackwell).

DANCY, Jonathan (1995), "Arguments from Illusion", *Philosophical Quarterly,* 45.

DARWALL, Stephen, GIBBARD, Allen and RAILTON, Peter (1992), "Towards *Fin de Siècle* Ethics: Some Trends", *Philosophical Review*, 101.

DAVIDSON, Donald (1982a), *Essays on Actions and Events* (Oxford, Clarendon Press).

DAVIDSON, Donald (1982b), "Expressing Evaluations" (Kansas, The Lindley Lecture, University of Kansas).

DAVIDSON, Donald (1984), *Inquiries into Truth and Interpretation* (Oxford, Clarendon Press).

DAVIDSON, Donald (1985a), "A Coherence Theory of Truth and Knowledge", in LePore, editor, *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson.*

DAVIDSON, Donald (1985b), "Empirical Content", in LePore, editor, *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson.*

DAVIDSON, Donald (1989a), "The Myth of the Subjective", in *Relativism: Interpretation and Confrontation*, M. Krauz, editor (Notre Dame, University of Notre Dame Press, 1989).

DAVIDSON, Donald (1989b), "Meaning, Truth and Evidence", in *Perspectives on Quine*, Robert Gibson and Roger Barrett, editors (Oxford, Blackwell, 1989).

DAVIDSON, Donald (1990), "The Structure and Content of Truth", *Journal of Philosophy*, 87.

DIVERS, John and MILLER, Alexander (1994), "Why the expressivist about value should not love minimalism about truth", *Analysis*, 54.

ENGBERG-PEDERSEN, Troels (1995), "Pitfalls in doing post-Hegelian ethics with Aristotle: a reply to Lovibond", in Heinaman, editor, *Aristotle and Moral Realism*.

EVANS, Gareth (1981), "Semantic Theory and Tacit Knowledge", in his *Collected Papers* (Oxford, Clarendon Press, 1985), and originally in Holtzman and Leich, editors, *Wittgenstein: To Follow a Rule* (page references to the reprinted version).

EVERSON, Stephen (1995), "Aristotle and the explanation of evaluation: a reply to David Charles", in Heinaman, editor, *Aristotle and Moral Realism.*

FOOT, Philippa (1972), "Morality as a System of Hypothetical Imperatives", *Philosophical Review*, 81.

FRENCH, Peter A., UEHLING, Theodore E. and WETTSTEIN, Howard K. (1992), *Midwest Studies in Philosophy Volume XVII: The Wittgenstein Legacy* (Notre Dame, University of Notre Dame Press).

GALLIE, W.B. (1955-6), "Essentially Contested Concepts", *Proceedings of the Aristotelian Society*, 56.

GEACH, Peter (1965), "Assertion", *Philosophical Review*, 74.

GIBBARD, Allan (1990), *Wise Choices, Apt Feelings* (Cambridge MA, Harvard University Press).

GOLDFARB, Warren (1985), "Kripke on Wittgenstein on Rules", *Journal of Philosophy*, 82.

HARE, Richard (1952), *The Language of Morals* (Oxford, Oxford University Press).

HARE, Richard (1963), *Freedom and Reason* (Oxford, Clarendon Press).

HARE, Richard (1981), *Moral Thinking: Its Levels, Method and Point* (Oxford, Clarendon Press).

HARMAN, Gilbert (1977), *The Nature of Morality* (Oxford, Oxford University Press).

HEGEL, G.W.F. (1942), *Philosophy of Right*, T.M. Knox, translator (Oxford, Clarendon Press).

HEINAMAN, Robert (1995), editor, *Aristotle and Moral Realism* (London, UCL Press).

HERMAN, Barbara (1993), *The Practice of Moral Judgment* (Cambridge MA, Harvard University Press).

HOLTZMAN, Steven H. and LEICH, Christopher M. (1981), editors, *Wittgenstein: To Follow a Rule* (London, Routledge & Kegan Paul).

HONDERICH, Ted (1985), editor, *Morality and Objectivity* (London, Routledge & Kegan Paul).

HOOKER, Brad (1996), editor, *Truth in Ethics* (Oxford, Blackwell).

HORWICH, Paul (1990), *Truth* (Oxford, Blackwell).

HORWICH, Paul (1994), "The essence of expressivism", *Analysis*, 54.

HUME, David (1975), *Enquiries concerning Human Understanding and concerning the Principles of Morals*, L.A. Selby-Bigge and P.H. Nidditch, editors (Oxford, Oxford University Press).

HUME, David (1978), *A Treatise of Human Nature*, L.A. Selby-Bigge and P.H. Nidditch, editors (Oxford, Clarendon Press).

HURLEY, S.L. (1984), "Frege, Proliferation of Force and Non-Cognitivism", *Mind*, 93.

HURLEY, S.L. (1985), "Objectivity and Disagreement", in Honderich, editor, *Morality and Objectivity.*

HURLEY, S.L. (1989), *Natural Reasons* (Oxford, Clarendon Press).

HURLEY, S.L. (1992), "Intelligibility, Imperialism and Conceptual Scheme", in French, Uehling and Wettstein, editors, *Midwest Studies in Philosophy Volume XVII: The Wittgenstein Legacy.*

JOHNSON, Paul (1989), *Wittgenstein and Moral Philosophy* (London, Routledge & Kegan Paul).

JOHNSTON, Mark (1989), "Dispositional Theories of Value", *Proceedings of the Aristotelian Society*, 63 (supplementary volume).

KANT, Immanuel (1929), *Immanuel Kant's Critique of Pure Reason*, Norman Kemp Smith, translator (London, Macmillan).

KANT, Immanuel (1988), *Foundations of the Metaphysics of Morals*, in *Kant: Selections*, Lewis White Beck, editor and translator (New York, Macmillan, 1988).

KOLNAI, Aurel (1969-70), "Moral Consensus", *Proceedings of the Aristotelian Society*, 70.

KORSGAARD, Christine (1996a), *Creating the Kingdom of Ends* (Cambridge, Cambridge University Press).

KORSGAARD, Christine (1996b), *The Sources of Normativity* (Cambridge, Cambridge University Press).

KRIPKE, Saul A. (1982), *Wittgenstein on Rules and Private Language* (Oxford, Blackwell).

LEAR, Jonathan (1982), "Leaving the World Alone", *Journal of Philosophy*, 79.

LEAR, Jonathan (1984), "The Disappearing 'We'", *Proceedings of the Aristotelian Society*, 58 (supplementary volume).

LePORE, Ernest (1985), editor, *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson* (Oxford, Blackwell).

LEWIS, David, (1989), "Dispositional Theories of Value", *Proceedings of the Aristotelian Society*, 63 (supplementary volume).

LOVIBOND, Sabina (1983), *Realism and Imagination in Ethics* (Oxford, Blackwell).

LOVIBOND, Sabina, and WILLIAMS, S.G. (1996), editors, *Essays for David Wiggins: Identity, Truth and Value* (Oxford, Blackwell/Aristotelian Society).

MACKIE, J.L. (1977), *Ethics: Inventing Right and Wrong* (Harmondsworth, Penguin).

McDOWELL, John (1977), "On the Sense and Reference of a Proper Name", *Mind*, 86.

McDOWELL, John (1978), "Are Moral Requirements Hypothetical Imperatives?", *Proceedings of the Aristotelian Society,* 52 (supplementary volume). Reprinted in McDowell (1998).

McDOWELL, John (1979), "Virtue and Reason", *The Monist*, 62. Reprinted in McDowell (1998).

McDOWELL, John (1980), "The Role of Eudaimonia in Aristotle's Ethics", in *Essays on Aristotle's Ethics*, Amelie Oksenberg Rorty, editor (Berkeley and Los Angeles, University of California Press, 1980). Reprinted in McDowell (1998).

McDOWELL, John (1981a), "Non-Cognitivism and Rule-Following", in Holtzman and Leich, editors, *Wittgenstein: To Follow a Rule*. Reprinted in McDowell (1998).

McDOWELL, John (1981b), "Anti-Realism and the Epistemology of Understanding", in *Meaning and Understanding*, Herman Parret and Jacques Bouveresse, editors (Berlin and New York, De Gruyter, 1981).

McDOWELL, John (1982a), "Criteria, Defeasibility and Knowledge", *Proceedings of the British Academy*, 68.

McDOWELL, John (1982b), "Reason and Action", *Philosophical Investigations*, 5.

McDOWELL, John (1983), "Aesthetic Value, Objectivity and the Fabric of the World", in *Pleasure, Preference and Value: Studies in Philosophical Aesthetics*, Eva Schaper, editor (Cambridge, Cambridge University Press, 1983). Reprinted in McDowell (1998).

McDOWELL, John (1984), "Wittgenstein on Following a Rule", *Synthese*, 58. Reprinted in McDowell (1998).

McDOWELL, John (1985a), "Functionalism and Anomalous Monism", in *Actions and Events: Perspectives on the Philosophy of Donald Davidson,* Ernest LePore and Brian McLaughlin, editors (Oxford, Blackwell, 1985). Reprinted in McDowell (1998).

McDOWELL, John (1985b), "Values and Secondary Qualities", in Honderich, editor, *Morality and Objectivity.* Reprinted in McDowell (1998).

McDOWELL, John (1986), "Critical Notice: *Ethics and the Limits of Philosophy* by Bernard Williams", *Mind*, 95.

McDOWELL, John (1987), "In Defence of Modesty", in *Michael Dummett: Contributions to Philosophy*, Barry Taylor, editor (Dordrecht, Martinus Nijhoff, 1987).

McDOWELL, John (1989), "One Strand in the Private Language Argument", *Grazer Philosphische Studien*, 33/34. Reprinted in McDowell (1998).

McDOWELL, John (1991), "Intentionality and Interiority in Wittgenstein", in *Meaning Scepticism*, Klaus Puhl, editor (Berlin and New York, De Gruyter, 1991). Reprinted in McDowell (1998).

McDOWELL, John (1992), "Meaning and Intentionality in Wittgenstein's Later Philosophy", in French, Uehling and Wettstein, editors, *Midwest Studies in Philosophy Volume XVII: The Wittgenstein Legacy.* Reprinted in McDowell (1998).

McDOWELL, John (1994), *Mind and World* (Cambridge MA, Harvard University Press).

McDOWELL, John (1995a), "Eudaimonism and Realism in Aristotle's Ethics", in Heinaman, editor, *Aristotle and Moral Realism.*

McDOWELL, John (1995b), "Might there be external reasons?", in *World, Mind and Ethics*, J.E.J. Altham and Ross Harrison, editors (Cambridge, Cambridge University Press, 1995).

McDOWELL, John (1995c), "Two Sorts of Naturalism", in *Virtues and Reasons: Philippa Foot and Moral Theory*, Rosalind Hursthouse, Gavin Lawrence and Warren Quinn, editors (Oxford, Clarendon Press, 1995). Reprinted in McDowell (1998).

McDOWELL, John (1998), *Mind, Value and Reality* (Cambridge MA, Harvard University Press).

McGINN, Colin (1983), *The Subjective View* (Oxford, Clarendon Press).

McGINN, Colin, (1984), *Wittgenstein on Meaning* (Oxford, Blackwell/Aristotelian Society).

McNAUGHTON, David (1988), *Moral Vision* (Oxford, Blackwell)

MOORE, A.W. (1996), "On There Being Nothing Else to Think, or Want, or Do", in Lovibond and Williams, editors, *Essays for David Wiggins: Identity, Truth and Value.*

NAGEL, Thomas (1970), *The Possibility of Altruism* (Oxford, Clarendon Press).

NAGEL, Thomas (1991), *Equality and Partiality* (Oxford, Oxford University Press).

NELL (O'Neill), Onora (1975), *Acting on Principle: an essay on Kantian ethics* (New York, Columbia University Press).

O'NEILL, Onora (1989), *Constructions of Reason: Explorations of Kant's Practical Philosophy* (Cambridge, Cambridge University Press).

PEACOCKE, Christopher (1981), "Reply: Rule-Following: The Nature of Wittgenstein's Arguments", in Holtzman and Leich, editors, *Wittgenstein: To Follow a Rule*.

PEACOCKE, Christopher (1984), "Review of *Wittgenstein on Rules and Private Language* by Saul A. Kripke", *Philosophical Review*, 93.

PEARS, David (1988), *The False Prison*, Volume Two (Oxford, Clarendon Press).

PLATTS, Mark (1979), *The Ways of Meaning* (London, Routledge & Kegan Paul).

PLATTS, Mark (1981), "Moral reality and the ends of desire", in *Reference, Truth and Reality*, Mark Platts, editor (London, Routledge & Kegan Paul, 1981).

PUTNAM, Hilary (1981), *Reason, Truth and History* (Cambridge, Cambridge University Press).

PUTNAM, Hilary (1996), "On Wittgenstein's Philosophy of Mathematics", *Proceedings of the Aristotelian Society*, 70 (supplementary volume).

RAWLS, John (1955), "Two Concepts of Rules", *Philosophical Review*, 64.

RAWLS, John (1971), *A Theory of Justice* (Oxford, Oxford University Press).

RAWLS, John (1980), "Kantian Constructivism in Moral Theory", *Journal of Philosophy*, 77.

RAWLS, John (1985), "Justice as Fairness: Political Not Metaphysical", *Philosophy and Public Affairs*, 14.

SCANLON, T.M. (1982), "Contractualism and Utilitarianism", in *Utilitarianism and Beyond*, Amartya Sen and Bernard Williams, editors (Cambridge, Cambridge University Press, 1982).

SEABRIGHT, Paul (1987), "Explaining Cultural Divergence: A Wittgensteinian Paradox", *Journal of Philosophy*, 84.

SIBLEY, F.N. (1968), "Objectivity and Aesthetics", *Proceedings of the Aristotelian Society*, 63 (supplementary volume).

SIDGWICK, Henry (1981), *The Methods of Ethics* (Seventh Edition) (Indianapolis, Hackett).

SMITH, Michael (1989), "Dispositional Theories of Value", *Proceedings of the Aristotelian Society*, 63 (supplementary volume).

SMITH, Michael (1994a), *The Moral Problem* (Oxford, Blackwell).

SMITH, Michael (1994b), "Why the expressivist about value should love minimalism about truth", *Analysis*, 54.

SMITH, Michael (1994c), "Minimalism, truth-aptitude and belief", *Analysis*, 54.

SNOWDON, Paul (1980-1), "Perception, Vision and Causation", *Proceedings of the Aristotelian Society,* 81.

STROUD, Barry (1984), "The Allure of Idealism", *Proceedings of the Aristotelian Society*, 58 (supplementary volume).

TANNER, Michael (1968), "Objectivity and Aesthetics", *Proceedings of the Aristotelian Society*, 63 (supplementary volume).

TAYLOR, Charles (1979), *Hegel and Modern Society* (Cambridge, Cambridge University Press).

WIGGINS, David (1980), "What Would be a Substantial Theory of Truth?", in *Philosophical Subjects*, Zak van Straaten, editor (Oxford, Oxford University Press, 1980).

WIGGINS, David (1990-1), "Moral Cognitivism, Moral Relativism and Motivating Moral Beliefs", *Proceedings of the Aristotelian Society*, 91.

WIGGINS, David (1995), "Eudaimonism and realism in Aristotle's ethics: a reply to John McDowell", in Heinaman, editor, *Aristotle and Moral Realism.*

WIGGINS, David (1996a), "Objective and Subjective in Ethics, with Two Postscripts about Truth", in Hooker, editor, *Truth in Ethics.*

WIGGINS, David (1996b), "Replies", in Lovibond and Williams, editors, *Essays for David Wiggins: Identity, Truth and Value.*

WIGGINS, David (1998), *Needs, Values, Truth* (Third Edition) (Oxford, Oxford University Press).

WILLIAMS, Bernard (1981), *Moral Luck* (Cambridge, Cambridge University Press).

WILLIAMS, Bernard (1985), *Ethics and the Limits of Philosophy* (London, Fontana Press).

WILLIAMS, Bernard (1996), "Truth in Ethics", in Hooker, editor, *Truth in Ethics.*

WILLIAMSON, Timothy (1994), "Critical Notice: *Truth and Objectivity*, by Crispin Wright", *International Journal of Philosophical Studies*, 2.

WINCH, Peter (1965), "The Universalizability of Moral Judgements", *The Monist*, 48.

WINCH, Peter (1981), "Im Anfang war die Tat", in *Perspectives on the Philosophy of Wittgenstein*, Irving Block, editor (Oxford, Blackwell, 1981).

WINCH, Peter (1991), "Certainty and Authority", in *Wittgenstein Centenary Essays*, A. Phillips Griffiths, editor (Cambridge, Cambridge University Press, 1991).

WITTGENSTEIN, Ludwig (1965), "A Lecture on Ethics", *Philosophical Review*, 74, and reprinted in *Philosophical Occasions 1912-1951*, James Klagge and Alfred Nordmann, editors (Indianapolis and Cambridge, Hackett, 1993).

WITTGENSTEIN, Ludwig (1966), *Lectures and Conversations on Aesthetics, Psychology and Religious Belief*, Cyril Barrett, editor (Oxford, Blackwell).

WITTGENSTEIN, Ludwig (1967), *Zettel*, G.E.M. Anscombe and G.H. von Wright, editors, G.E.M. Anscombe, translator (Oxford, Blackwell).

WITTGENSTEIN, Ludwig (1975), *The Blue and Brown Books*, R. Rhees, editor (Oxford, Blackwell).

WITTGENSTEIN, Ludwig (1976), *Philosophical Investigations*, G.E.M. Anscombe, R. Rhees and G.H. von Wright, editors, G.E.M. Anscombe, translator (Oxford, Blackwell).

WITTGENSTEIN, Ludwig (1977), *On Certainty,* G.E.M. Anscombe and G.H. von Wright, editors, Denis Paul and G.E.M. Anscombe, translators (Oxford, Blackwell).

WITTGENSTEIN, Ludwig (1978a), *Culture and Value* (Second Edition), G.H. von Wright, editor, Peter Winch, translator (Oxford, Blackwell).

WITTGENSTEIN, Ludwig (1978b), *Remarks on the Foundations of Mathematics* (Third Edition), G.H. von Wright, R. Rhees, G.E.M. Anscombe, editors, G.E.M. Anscombe, translator (Oxford, Blackwell).

WITTGENSTEIN, Ludwig (1993), "Remarks on Fraser's Golden Bough", in *Philosophical Occasions 1912-1951*, James Klagge and Alfred Nordmann, editors (Indianapolis and Cambridge, Hackett, 1993).

WRIGHT, Crispin (1980), *Wittgenstein on the Philosophy of Mathematics* (London, Duckworth).

WRIGHT, Crispin (1984), "Kripke's Account of the Argument Against Private Language", *Journal of Philosophy,* 81.

WRIGHT, Crispin (1986), "Does *PI* 258-260 Suggest a Cogent Argument Against the Possibility of a Private Language?", in *Subject, Thought and Context*, Philip Pettit and John McDowell, editors (Oxford, Clarendon Press, 1986).

WRIGHT, Crispin (1988a), "Moral Values, Projection and Secondary Qualities", *Proceedings of the Aristotelian Society,* 62 (supplementary volume).

WRIGHT, Crispin (1988b), "Realism, Anti-realism, Irrealism, Quasi-realism", in *Midwest Studies in Philosophy Volume XII: Realism*, Peter A. French, Theodore E. Uehling and Howard K. Wettstein, editors (Minneapolis, University of Minnesota Press, 1988).

WRIGHT, Crispin (1992), *Truth and Objectivity* (Cambridge MA, Harvard University Press).

WRIGHT, Crispin (1993), *Realism, Meaning and Truth* (Second Edition) (Oxford, Blackwell).

WRIGHT, Crispin (1994), "Realism, Pure and Simple? A Reply to Timothy Williamson", *International Journal of Philosophical Studies*, 2.

WRIGHT, Crispin (1996), "Truth in Ethics", in Hooker, editor, *Truth in Ethics.*