OBJECTIVITY
AND
MORAL JUDGEMENTS

OBJECTIVITY
AND
MORAL JUDGEMENTS

A Thesis
Presented for the Degree of
Doctor of Philosophy in Philosophy
by

CHRISTOPHER WILLIAM BOYNE

University of Kent at Canterbury

August 1983

I would like to acknowledge my debt to a good many people who have helped me in various ways in the writing of this thesis. I am grateful to Professor D.W.Hamlyn, for first suggesting to me that some aspect of objectivity would be a worthwhile research topic, to Professor R.J. Butler, who read part of an early draft and made valuable suggestions, and to Mr Robin Taylor for a discussion of some points on Kant. My greatest debt, however, is to my supervisor, Dr Colin Radford, not only for his guidance in academic matters but also for his unfailing encouragement and support. Support, encouragement and tolerance have also been forthcoming from my colleagues at South Kent College of Technology, from many friends, and most of all from my mother, who had the most to tolerate.

My sincerest thanks go to them all.

CONTENTS

ABSTRACT OF THE THESIS

The assumption that moral judgements could never be objective stems from a series of misunderstandings about objectivity itself. Objectivity is misleadingly identified with true judgements, or with judgements about material objects, and seen as exclusive to empirical judgements. But what makes an empirical judgement objective is itself hardly better understood. In particular, the legacy of empiricism, the belief that the justification for the objectivity of empirical judgements must lie within individual experience, is inadequate. And the account which must replace it is one that makes our normal agreement about the conclusions of our fundamental judgements a precondition of their intelligibility. This account is applicable not only to the field of empirical judgements, but to any area of judgement where questions of justification can arise. The assumption that justification can only be found in factual, never in evaluative, judgements is challenged, and I argue that our ability to give reasons for moral judgements presupposes that we normally agree in the basic judgements we make about what has moral value. The belief that moral judgements could be made from outside the standpoint created by this normal agreement is false; such judgements would be unintelligible. Thus morality necessarily involves the existence of a framework of fundamental value-judgements which are objective, though these judgements have only a limited capacity to determine how we ought to behave.

CHAPTER ONE

SOME MISCONCEPTIONS ABOUT OBJECTIVITY

The question of whether moral judgements can ever be objective has exercised the minds of many of the ablest philosophers since the time of Plato. Nevertheless, in recent years there has been a tendency to assume that the question is not really a difficult one at all. Sometimes it has been suggested that the question cannot even be sensibly asked; more commonly it has been thought that although we can ask the question easily enough, it can also very easily be answered, and in the negative. I believe, however, and in this chapter I shall try to show, that the question of moral objectivity is being dismissed rather too lightly. Easy answers often derive from a failure to appreciate what is involved in the question they answer, and here the usual answers have their roots in a number of considerable confusions about what moral objectivity amounts to. This is less a matter of confusion about the nature of moral thought (though elements of that are certainly involved) than a matter of confusion about the concept of objectivity itself. So although the central topic of this thesis is a question belonging to moral philosophy, we shall be concerned in a good many of the pages that follow with getting clear about the implications of the term 'objective judgement' in its more familiar empirical context. I make no apology for this, since philosophy of all disciplines should resist the temptation to become too narrow and compartmentalised in its approach to any problem; but I hope that this chapter will show why we need to approach the topic in this way, by indicating ways in which the concept of objectivity has been confused with quite different notions, and suggesting that we need to reach a clearer understanding of the concept itself before we can profitably consider whether it ever applies to moral judgements.

But first of all we ought to consider why the question of moral objectivity should even arise. Why should it occur to us to wonder whether moral judgements might ever be objective? The answer that is most likely to suggest itself to us initially is the least satisfactory one, because it is based on psychological rather than philosophical considerations. It is the idea that moral beliefs are too important to us to be thought of as no more than personal views or socially determined convictions. We wish to believe that our moral standards have some sort of objective validity, either because we fear the thought of being responsible for creating them ourselves, or because we desire a greater basis of certainty for them than we feel our own belief, or the belief provided by the society we live in, can provide. What we would like to find, however, is no guide to what we actually will find in an enquiry which should itself attempt to be objective; and the psychological urge to find objectivity in moral matters, though a dominant trait in the history of human thought, has no place in a philosophical study.

Psychological promptings apart, we may feel philosophical doubts about the accounts given by many modern moral philosophers of the nature of moral thought. They draw a contrast between factual and moral judgements in terms of the contrast between the intellect and the will, and characterise moral thinking in terms of our deciding on moral principles rather than on discovering them. However, as Bernard Williams puts it, in what he describes as a "gesture towards a centre of dissatisfaction", "the consciousness of a principle of action as freely decided upon is very unlike the consciousness of a moral principle, which is rather of something that has to be acknowledged. If it is then said that there is just a psychological explanation of that, then moral thought seems a cheat, presenting itself to us as too like something which it is not."* The idea that moral thought looks as if it is constrained towards some conclusions rather than others is scarcely enough on its own to entail

* B.Williams 'Morality' CUP 1976, p.50

the objectivity of moral principles, as Williams himself later points out (ibid.). But this does indicate the possibility of raising the question of moral objectivity as a genuine philosophical issue, which a psychological explanation would merely mask and not settle. For why should moral judgement give this impression of being independent of the promptings of the individual will?

Another reason for considering the possibility of moral objectivity occurring in some form or other is the fact that moral judgements do not appear merely to describe or reflect personal beliefs or attitudes. It is not just that 'it is good' seems to mean something different from 'I like it', but that the former, unlike the latter, seems to refer to something independent of anyone's beliefs or attitudes. G.E.Moore remarks on this in Chapter 3 of his 'Ethics', in a discussion of moral disagreements.* He argues that if the rightness or wrongness of an action were to depend on any person's, or group of persons', feelings or thoughts, and if one person or group thought a particular action was right while another person or group thought it was wrong, the same action would be both right and wrong. The consequence that Moore derives from this, the idea that if this were so moral disagreement would be impossible, is true only if moral judgements are taken to be nothing other than reports of personal attitudes; and more recent analysis of the logic of moral discourse shows them to be something rather different. Nevertheless, Moore has a point in wondering why the distinction between being right and seeming right to someone is applicable to moral judgements, when it is not applicable to statements about personal tastes or preferences. This can be illustrated if we consider what someone would say of an action about which he had changed his moral opinion over a period of time. He might well say that at one time he believed the action to be right, but now he can see that it was wrong. But if moral judgements depend only on personal beliefs, and do not reflect something independent of those beliefs, what are the words 'believed' and 'see' doing in the previous sentence? Isn't it more natural to say that the action

* OUP 1912

was right but is now wrong? Why do these words give the impression of there being something independent of the individual's beliefs here? Is it just that moral language apes language about objects; or is it just that the universalising character of moral language forces us to speak like this? And why in any case should moral language have this universalising character if what it expresses is only a personal viewpoint? Again, the possibility of moral objectivity is one of the range of answers which deserves consideration.

Finally, there is an aspect of the question which has special relevance for parents and teachers because it is they in particular who have to give moral instruction and guidance to the young. Sometimes moral instruction can take on a relatively neutral guise, as when one merely points out the consequences of different courses of action; but moral guidance can rarely if ever be entirely neutral, and arguably it should not be anyway. Since in the roles of parent or teacher we transmit our beliefs to others in a situation where our beliefs carry particular weight, the status of what we are saying is important, both for what we say and for how we say it. Reflecting on this may well force us to consider whether what we are doing is exercising a personal influence, or acting as agents for the transmission of our society's mores, or again conveying something which has points of similarity with other branches of knowledge. So in this way too, the question of whether and to what extent our moral beliefs could count as objective can arise as an issue.

Now against these slightly tenuous considerations, modern philosophers have ranged a whole battery of arguments, of varying quality and sometimes mutual incompatibility, to show that moral judgements could not possibly be objective. It is argued that the bewildering variety of actual moral responses that occur, both within societies and among different ones, demonstrates the impossibility of moral objectivity. It is said that historical changes in moral beliefs show the same thing. It is suggested that moral objectivity is incompatible with free will and with the notion of individual moral

responsibility. We are told that moral judgements could not be objective since nothing could count as a moral object, and that the requisite faculty of moral perception does not exist. We are told that moral judgements could not be objective since moral judgements are not statements, and because moral judgements are not factual judgements, and finally because moral judgements are neither true nor false. I shall not discuss all these objections here and now, but allow them to wither away gradually in the general course of my argument. For there is really no advantage in trying to meet all these objections whilst we are still so unclear about the concept that is central to them all. As we become clearer about what is really meant by calling a judgement objective, it will be evident that in their different ways they all miss the point.

The fact that there is confusion about the meaning of objectivity has been quite widely recognised. D.W.Hamlyn remarks, for instance, that "the notion of objectivity has been often misunderstood,"* and Professor Körner has written "Philosophical reflection often leads to the question whether an entity is real or a principle objective. Because of an almost hopeless instability in the use of the term 'objective' the question is often confused and confusing. 'Objectivity' is used in completely different senses and often with no clear meaning at all."** (It might be added that Körner faithfully reflects the instability of usage he refers to. On the same page as these words he also writes "the truth (the objectivity) of any of the principles" and "the objectivity or, if we like, the absolute validity of the categorical imperative". The second phrase seems to imply, as does Kant's own usage, that the objectivity of a principle involves its necessary truth. If so, Kant was using the term 'objectivity' in a special and technical sense that it does not usually bear. But because of his immense influence on subsequent philosophy, elements of this technical usage spill over into other areas of philosophy

* 'Objectivity' in 'Education and the Development of Reason' ed. Dearden, Hirst and Peters, Routledge, 1972

** 'Kant', Penguin Books 1955, p.142

and add to the instability Körner complains of. Since I am not engaged in a study of the history of the idea of objectivity, and because I do not find Kant's special use of the term helpful in my discussion of it, I shall not discuss what Kant meant here. But because it would be wrong simply to ignore Kant on this matter, I have located some comments on Kant and objectivity in an appendix.)

The instability of usage referred to above is one reason why the terms 'objective' and 'objectivity' do not occur very often in modern Ethics, but it is not the only reason. Another is that the term was annexed (or is widely thought to have been) by Ethical Intuitionists at the beginning of this century, so that often a reference to moral objectivity is taken to be a reference to the tenets of a discredited moral theory.* Faced with all this confusion, it is not easy to see where clarification of the concept could best begin. So I will start at something of a tangent by looking at some fairly recent references to objectivity in the context of evaluative judgement, in order to see what one or two philosophers have thought objectivity amounted to. Apart from being the very first stage of an analysis of the concept, this will also serve to show us why the possibility of there being objective moral judgements is so readily written off by many philosophers.

The concept of objectivity is sometimes referred to in Ethics as part of the process of drawing a contrast between factual and non-factual judgements. This is usually a preliminary to drawing a contrast between factual judgements and <u>moral</u> judgements, and may predispose the reader to identify moral judgements with other types of non-factual judgement in a slightly underhand way. Unfortunately, the reference to objectivity only aids this. However, the argument is so familiar to the experienced philosopher, and seems so straigntfoward to the novice, that it has not attracted the critical scrutiny it deserves. I quote an example of this sort,

*see e.g. R.M.Hare 'The Language of Morals', OUP, p.77

taken from an introductory book aimed at the general reader and the pre-university student:

"In this chapter we are going to try to analyse and examine comparisons in general, and in particular the whole set of comparisons or implicit comparisons which are called value judgements...We start by considering two examples, one at each end of a rather important scale.

'Objective' Comparisons

I have in front of me two sticks, and I say that this one (A) is longer than that one (B). It would be generally agreed that whether this is true or not is a matter of fact. The statement can be subjected to a public test by putting them beside one another, and, supposing for the moment that the difference is obvious to the naked eye and that we are not entertaining eccentric metaphysical notions about the nature of reality, there would be no difficulty in coming to an agreed decision about it...It might happen, however, that to the naked eye the two sticks appeared to be about the same length. One person might then say he thinks A is longer, while someone else thinks B is. A scientist with instruments for measuring accurately might then be called in to decide between them. Although the two observers heve formed different views about the comparative lengths of the sticks, they would be likely to agree that it is nevertheless true that which is the longer is a matter of fact and not of opinion, that one of them is right and the other wrong...though the facts may be difficult to discover...This comparison of length is a matter of fact which may be subjected to a public, scientific test. Any statement that is made about it can be, at least in principle, verified or falsified...Such a statement...that can be verified or falsified by experience (empirically), is often called <u>objective</u>. It is easy to think of whole classes of comparisons of a similar kind...

'Subjective' Comparisons

We come now to our second example. Suppose that someone tries out two chairs by sitting in each of them one after the other and then says: 'This chair is more comfortable than that'. It is probable that if questioned he would

agree that what he really means is that he personally finds the one chair more comfortable than the other, with perhaps the implication that most other people are likely to do so too. In other words, although the statement may appear at first to be of the same kind as 'this chair is heavier than that', it becomes clear as soon as we think about it that the speaker is likely to be expressing or perhaps describing his own attitude to the chair...Statements which express the attitude of the speaker are often called subjective...

Both the cases that we have considered so far are obvious and easy, and it may seem surprising that it has been thought worth while to discuss them at such length. Our reason for doing so is that the consideration of these cases at the two extremes provides a solid base from which to work for the discussion of trickier borderline cases."*

The trickier borderline cases that the author has in mind are, of course, moral and aesthetic judgements, which he would admit are not quite equivalent to expressions of personal preference. But the lines of demarcation, within which discussion will take place, have already been set out. The alternatives on offer are objective judgements (like statements of fact) and subjective judgements (like personal attitudes). From here it is an exceedingly small step to assuming that only facts (and only certain sorts of facts at that) count as objective, and anything that fails to qualify as a fact is not going to be objective either.

The general lines of the argument I have quoted will, I imagine, be only too familiar. I have chosen this example of it to quote at length only because of the way its author gives prominence to the terms 'objective' and 'subjective' in the course of what he says. This is convenient because it is the ready assumption about the place of objectivity in the fact-value contrast that I particularly want to challenge. Most moral philosophers who draw the fact/non-fact distinction in the same way

* E.R.Emmet 'Learning to Philosophize', Penguin, pp.109-112

manage nevertheless to do so without mentioning objectivity as such. It is, as I have said, something of a taboo word; it is sometimes introduced at a later stage, always used sparingly, and invariably presented with little or no reference to what it is supposed to signify. Yet the assumption that, facts being objective, anything not a fact is thereby non-objective, often lurks persistently under the surface of their arguments.

Emmet, at least, comes out into the open a little more, and it is fairly clear what he takes 'objective' to mean. For him, an objective judgement is one which, if it were true, would state a fact (typically) about the external world. So much seems to be implied, at any rate, by his claim that statements are called objective when they can be empirically verified or falsified. Though he does not confine the application of the term 'objective' to what can be empirically tested, it is fairly clear that the terms 'factual', 'empirically testable' and 'objective' are being lined up against 'non-factual', 'personal' and 'subjective' in such a way that it becomes inevitable that anything non-factual which looks at all like an expression of personal opinion will be branded subjective. To be fair to Emmet, he was not writing a scholarly work, and he has not offered a formal definition of 'objective judgement', but only said that judgements purporting to state facts are "often called" objective. (It is for that reason that I used the word 'typically' above when I stated what I take Emmet's notion of objectivity to be.) But I am sure that it has simply not occurred to him that the fact/non-fact distinction needs to be separated from the objective/subjective distinction in any way. He has, after all, introduced the terms 'objective' and 'subjective' in order to label the distinction between what is a matter of fact, and therefore a matter where agreement is proper, and what is merely a matter of opinion, where agreement is neither required nor expected. So Emmet's way of approaching the fact-value distinction is untypical only in that he is a little less coy than most moral philosophers in suggesting that only factual judgements could count as objective.

Before I attempt to criticise the view that facts are typically what is objective, and anything which is not a fact is thus going to be subjective in the way expressions of personal preference are subjective, I would like to look briefly at the way in which this view is often developed. The words are my own, but they are based on a well-known summary of the views attributed to Moral Prescriptivism.* The argument is developed as follows: We have seen that we can speak of empirical judgements as being right or wrong. Thus, in principle, it is always possible to say that a person's judgements about the empirical world are correct or erroneous. On the other hand, moral judgements, thought they may take account of the facts, are never absolutely determined by them. For it is always possible that two people will agree on what the facts are, and that they are relevant to a particular moral issue, and yet will disagree completely in the conclusions that they draw from those facts. Even more radically, there is no reason why either of them need accept just those facts as the ones relevant to a moral decision in the case before them. One person may indeed find nothing requiring a moral response in circumstances which another person finds morally compelling. For in moral judgement there is neither an agreed procedure for deciding what constitutes relevance, nor about what follows from anything that <u>is</u> agreed to be relevant. So (the argument continues) whereas the correctness of empirical judgements is determined by what is the case, independent of <u>beliefs</u> about what is the case, moral judgements are not. Moral judgement is autonomous, which means that it is open to every man to judge in his own fashion. No man is bound to judge in a particular way because of the facts, or because of the way other people happen to view the facts. Thus moral judgements could only express beliefs, the terms 'correct judgement' and 'incorrect judgement' have no place in moral discourse, and moral judgements can neither be true nor objective.

On the face of it, this argument is a simple and

* see P.Foot, 'Moral Beliefs', PAS 1958

elegant statement of the relatively obvious, and is moreover apparently confirmed by what we all know about the great variety of moral views that people actually hold. It is not often noticed, then, that it contains a fallacy. The question of whether people can disagree, or whether they do disagree, about moral questions is quite beside the point when the point is whether moral judgements could be objective. Nobody disputes that people can and do disagree about matters of fact, and nobody suggests that this is a reason for claiming that factual judgements are not capable of being objective. Why, then, should the variety of moral responses be taken as proof in itself that moral judgements cannot be objective? It may instead prove that people are often not very good at making moral judgements correctly, just as people are not very good at describing what happened at an accident they witnessed; or there may be all sorts of other reasons why people do not agree in their moral judgements. The fact of this disagreement is no reason for saying that some of them are not at any rate being objective. And the same argument can be put forward in order to dispute the significance of the suggestion that there is widespread disagreement over questions of relevance of data for basing one's judgement on. I do not dispute that where disagreement is widespread and seemingly unresolvable, we are entitled to ask why this should be; but the answer is not automatically that here nothing counts as right or wrong. This would be correct only if it could be shown that people could not hold the same opinions as one another, except perhaps fortuitously; that is, by showing the impossibility of there being criteria of correctness or relevance in moral judgements. But all the argument does is to assert that disagreement shows the absence of criteria of these sorts. So it is false to suppose that this argument disproves the possibility of moral objectivity. For criteria for correct judgement could exist even if people regularly failed to apply them properly, or else varied in their judgements for completely different reasons.

The argument which I am criticising gets its plausibility by covertly reducing moral judgements to a status of dependence on the individual which is characteristic of expressions of personal preference or taste. For these are such that every man is his own authority, and the fact that one man has this particular preference is no reason why anyone else should have it too. (I do not mean to suggest that a person can never be mistaken about his own tastes, or that others can never know better than he does what his tastes really are. But if a man is mistaken here, he is mistaken in his beliefs about his tastes, not in the tastes themselves. It makes no sense to say that a man is mistaken in his tastes.) There was, as we saw, a tendency within Emmet's argument to identify the factual and the objective, and there is a corresponding tendency to identify the non-factual (and hence, supposedly, the non-objective) with expressions of personal preference. The assumption is that there are basically only two types of judgement, those warranted by verification and those warranted only by the preferences of the individual who makes them. The "trickier borderline cases" that Emmet is interested in are tricky only because, he thinks, we will have to decide whether such things as moral judgements are really attenuated facts or really attenuated preferences.* Only the factual will count as objective, he thinks, because then verification becomes possible.

But if we accept this rather crude bifurcation of judgements, an odd consequence arises. For real disagreement over anything other than what is empirically verifiable becomes impossible, as does real agreement. Now this is the case where the question is one of personal preferences. If, for example, one man says 'I like strawberries' and another says 'I don't', we would not regard them as being in disagreement over something, but merely as having and expressing different tastes. In the same way, if two men both say that they like strawberries, they are not, strictly speaking, in

* Emmet, op. cit. pp.119-129, esp. p.127

agreement. They just have a common liking for strawberries which they both express. We would not expect the lover of strawberries to attempt to refute the man who finds them distasteful. He may give reasons for his liking them, in the sense of pointing out, for instance, what particular features of strawberries one should pay attention to when forming an opinion of them. But here he can do no more than hope that what he says will encourage the other man to reconsider his <u>own</u> opinion of the fruit. Tastes can change, and palates can of course be educated. But tastes cannot <u>be</u> changed by one person showing another that he was in fact mistaken in his judgement. The point is neatly summed up by R.W.Beardsmore when he remarks: "No-one can teach me to prefer tobacco to ice-cream...because in this context there is no such thing as judgement."*

Now if the expression of opinions which involve genuine judgement <u>were</u> like the expression of personal tastes or preferences, it would indeed be impossible to avoid the conclusion that moral judgements and the like could not be objective. For it could then be shown that it was not possible for people either to agree or to disagree with each other in their opinions, since there would be nothing independent of their own personal beliefs for them to agree or disagree about, as I have just argued. However, the sorts of evaluative judgement which arguments like Emmet's are really aimed at are <u>not at all</u> like expressions of personal preference, and it is thoroughly misleading to imply that they are. Let us consider one or two of the consequences of that implication.

If all non-factual judgements had the same status as expressions of taste or preference, it would follow that there could never be grounds for accepting another's judgement instead of one's own; or, if faced with a choice between two views, there would never be good or sufficient grounds for an individual to accept one and

* 'Art and Morality', Macmillan 1971 p.39. I agree with Beardsmore about the need to recognise that 'judgement' ought to be confined to what is interpersonal. In speaking of moral judgements, I am of course following accepted terminology, and not trying to prejudge the issue!

reject the other. I should make it clear at this point that by 'grounds' I mean something founded in reason, such that a man would feel himself bound to accept a certain view, although he might prefer to accept a different one. I mean something different from what he would desire to believe, or from what it is convenient or easy to believe, or from what it is customary for members of his society to believe. Most societies have a good many conventional beliefs or attitudes, which may simply express the consensus of feeling within that society about the topic in question; but an appeal to such a convention would not in itself be what I mean by a ground for accepting a judgement.

Now it is very far from the case that if a judgement is non-factual there cannot be something at least approaching grounds in the above sense for accepting or rejecting it as the case may be. One area where this is clearly so is the field of expert judgement. No doubt it is true that an expert is someone who knows more or who can tell more about his subject than others can. Yet it is wrong to suppose that expert judgement could never involve more than a specially wide or deep knowledge of facts or a special capacity for observation. For many of the fields in which we acknowledge that there are experts are fields where the most important judgements are evaluative rather than purely factual. Notoriously, experts disagree with one another, whether the issues are predominantly factual or whether they are primarily evaluative. Yet this is not incompatible with their status <u>as experts</u>; and we do not think of the expert as one who merely expresses a personal judgement which is in no way different from anyone else's, when he is called upon to make an assessment rather than to discover or reveal a fact, except that <u>his</u> evaluation is backed by more knowledge than the average man can muster. (Indeed, if the contrast between factual and personal judgements which I am attacking here were correct, we could not think of the expert in quite this way anyway. We might treat his judgement as no less a matter of personal opinion than anyone else's. But we could not regard his special knowledge as anything that

had a bearing on his capacity to make evaluations. For ex hypothesi evaluative judgements need not take any particular fact into account.) Our view of the expert is of one who is not only particularly knowledgeable in his field but who also has a capacity for sound judgement. He must be able to recognise what is relevant to judgement and must be able to assess its significance, going as far as the evidence warrants and no further. It is because of these capacities that we regard him as an expert in the full sense, and regard his judgements as likely to be right, and so trustworthy. So although we often cannot submit the expert's judgement to any form of verification or falsification, we are nevertheless entitled to give his judgements more credence than we would those of the layman, just because he is the expert and they are not. Loosely speaking, then, the expert is one who, inter alia, has a special capacity for judging aright.

I do not mean to suggest that expertise is invested with any sort of magical aura of correctness. I have only drawn attention to the status of experts so as to show our belief that there are areas of non-factual judgement where one man's judgement is better than another's, and where we cannot claim that every man is his own authority. If that is so, it follows that there must be criteria of some sort for the worth or lack of worth of judgements which are incapable of being determined simply by reference to the standards of empirical truth or falsity. Such criteria might not turn out to be sufficient to make an expert's judgement objective; but we should not ignore this indication that non-factual judgements can be made on what people hold to be good grounds.

Yet if the characterisation of non-factual judgements given earlier was correct, we would have to ignore the possibility that some non-factual judgements are made on better grounds than others. For ex hypothesi, if every man could decide what was to count as relevant for himself, and what followed from what he decided was relevant, again for him, there could be no such thing as grounds for such a judgement, whether good or bad. For in matters where

every man is his own authority, there is no authority. Thus we must either abandon the suggestion that all non-factual judgements are personal and made without any basis other than that of personal predilection, or else we must deny that there could be any such thing as expert judgement in non-factual matters, those cases where we sometimes speak of expert opinion. Faced with this choice, our scepticism is likely to be directed against the unhappy fact/non-fact contrast which I have been discussing, as being a glaring over-simplification.

How inappropriate that contrast is to our ordinary view of moral judgement can be shown by the following consideration. The corollary of the idea that there can never be good or sufficient reasons for accepting another person's non-factual judgements is that there can never be good or sufficient reasons for rejecting another person's non-factual judgements either. But now suppose that we are confronted with someone who professes a very odd moral opinion for which he gives a very odd reason. For example, a man might hold that the moral worth of one's actions depends overridingly on the length of one's fingernails. Let us suppose that this man shows his understanding of moral terms by judging quite normally in all cases except those where a person involved has very long fingernails. Such a person's actions, he believes, are always to be condemned, no matter what he has done or not done. Now we would surely regard such judgements as these as being awry. It would not just be because we might happen to judge differently in many of the cases where the people involved had long fingernails, or that we chose to regard different features of those situations as what would count for us as the morally relevant considerations. We would, I am sure, regard his judgements as being faultily based, at least insofar as that reason entered into them. This would still be our view if it happened that his judgements were in complete accord with our own. In other words, we must take into account another person's reasons for judging as he does when estimating the worth of his judgement. Yet on the view that all non-factual judgements are to

be treated as personal judgements, that is just what we cannot do. The strongest objection that we could make to this man's moral outlook is to say that it is unconventional. This, surely, is far too weak a condemnation. It is in just this sort of case that we could say with assurance that the man's reasoning was irrelevant or irrational. Yet according to the view I am criticising, relevance and rationality have no place in such judgements, or in our assessment of such judgements. For considerations of relevance and rationality are supposed to have no place in moral judgement at all except to the extent that the individual decides that they shall have for him.

From what I have said so far, then, it should be clear that the innocent-looking contrast between what can and what cannot count as an objective judgement, drawn in the sort of way that Emmet has drawn it, for instance, is prone to lead to a view of objectivity which is at best over-simple, and at worst thoroughly misleading. To equate the objective with the putatively factual, and to contrast all other types of judgement to that, results in a highly distorted view of what non-factual judgements can be like. Of course, nothing I have said so far has been enough to show that there are any non-factual judgements which are indeed objective. What I have shown, however, is that a theory which effectively excludes the possibility of there being good grounds, or indeed anything which deserves to be called a ground of any sort, for making any judgement other than one which can be empirically verified, must be mistaken in its view of what objectivity amounts to. For it would be odd indeed if the concept of objectivity had nothing to do with having grounds for one's view of things, and equally odd if the only grounds for judgement were empirical verification. Certainly part of what is meant, in ordinary life, by calling someone's view objective, whatever it is a view of, is that there are good reasons for his seeing things in the way that he does, as opposed to his having strong personal feelings which make him see things in that way. But the identification of the objective with the factual and that alone, has effectively blinded many philosophers

to the possibility that there might be other grounds than verifiability which would suffice for objectivity, if these grounds were not excluded from consideration at the outset.

The supposed link between objectivity and empirical verification, which I have been trying to undermine, leads readily into a further confusion which attempts to link objectivity with judgements about objects. This is not surprising; for many sorts of objects qualify as objects just because their existence is empirically verifiable. And when we think of cases of empirical verification, we are most likely to think of those where the statement which is verified asserts that some property or other belongs to some material object. However, it is not to be assumed that the objects, in virtue of reference to which judgements are sometimes supposed to be objective, have to have the characteristics of material objects. The objects to which objective judgements in Mathematics and Ethics are sometimes thought to refer are rarely conceived as material objects, but they are usually conceived on an analogy with material objects, usually because they are supposed to exhibit the same independence from the beliefs of the individual as material objects, and because the way in which one is supposed to come to know them is often treated as being akin to the process of knowing about a material object by inspection or observation. It is for both these reasons that the Ethical Intuitionists were sometimes known by the alternative name of Objectivists.* Ethical Intuitionists conceived of moral judgements as statements of quasi-fact about some realm to which the mind of the individual has direct access by inspection through what was often thought of as a special faculty rather like a faculty of sense. And objections to Intuitionism are usually, and quite rightly, attacks on the idea that this kind of knowing could be anything like the sort of knowing which involves the ordinary senses and ordinary (material) objects. I shall not discuss these objections at any length; they are too well known to need exposition, and in any case

* see Hare, loc.cit.

the difficulty in Intuitionism which they expose will appear in due course in another and more general form. In essence, these criticisms turn on the impossibility of drawing a distinction between knowledge of an object and beliefs about an object, or supposed object, when the 'object' in question is only available to the mind of the individual, and to no-one else's. This is a difficulty which is also relevant to the question of how the naive realist knows of the existence and the qualities of material objects. I shall return to this question towards the end of this chapter, and again in subsequent chapters. The interested reader will find the conventional objections to Ethical Intuitionism well set out in W.D.Hudson's 'Ethical Intuitionism' and P.H.Nowell-Smith's 'Ethics'.*

In the 'Blue Book', Wittgenstein wrote: "Philosophers constantly see the method of science before their eyes, and are irresistably tempted to ask and answer questions in the way science does. This tendency is the real source of metaphysics and leads the philosopher into complete darkness."** Treating 'the method of science' broadly, so as to include the organised process of determining independent existence through observation and verification of any sort, then Wittgenstein's remark is wholly apposite to the later forms of Ethical Intuitionism, where moral terms such as 'good' were treated as the names of non-natural properties of states of affairs, and moral judgements were treated as observation-statements which predicated those properties. G.E.Moore in particular rejected what was surely the bedrock of earlier Intuitionist belief, the notion that the moral faculty is unerring in its judgement or apprehension, when he wrote: "Still less do I imply (as most Intuitionists have done) that any proposition whatever is true, <u>because</u> we cognise it in a particular way or by the exercise of any particular faculty: I hold, on the contrary, that in every way in which it is possible to cognise a true proposition, it is

* Macmillan 1967 and Penguin Books 1954 respectively
** The Blue and Brown Books, Blackwell 1958, p.18

also possible to cognise a false one."* Here Moore was, I think, attempting to put claims to moral knowledge on exactly the same footing as claims to empirical knowledge by observation. This is not merely analogous with the ordinary process of gaining knowledge of the external world by observation, but actually an attempt to align moral introspection with scientific observation, science being the most fruitful source of accounts of the independent existence of entities, and relatively untrammelled, as Moore's ethical theory tries to be, with metaphysical considerations about the status of its subject-matter.

This conscious or unconscious alignment of ethical method with scientific method survived the general abandonment of Ethical Intuitionism, reappearing in the form of arguments for rejecting the possibility of moral objectivity simply on the grounds that moral objectivity suggests the existence of moral objects of questionable ontological status. This impatience with the notion that there might be existent moral objects for moral judgements to be judgements about, is evident in the analogy suggested by A.J.Ayer in the following passage:

"The familiar subjective-objective antithesis is out of place in moral philosophy. The problem is not that the subjectivist denies that certain wild, or domesticated animals 'objective values' exist and the objectivist triumphantly produces them; or that the objectivist returns like an explorer with tales from the kingdom of values and the subjectivist says he is a liar. It does not matter what the explorer finds or does not find. For talking about values is not a matter of describing what may or may not be there, the problem being whether it really is there. There is no such problem. The moral problem is: What am I to do? What attitude am I to take? And moral judgements are directives in this sense."**

Here Ayer goes beyond the verificationist-factualist

* 'Principia Ethica', Cambridge 1903, p.x

** 'On the Analysis of Moral Judgements', reprinted in his 'Philosophical Essays', Macmillan 1963, p.242

view of objectivity which I have attributed to Emmet, and implies that questions about objectivity are in effect questions about whether something does or does not have an independent existence. He maintains that the question of objectivity is out of place in Ethics because questions of existence are out of place there. The reason he suggests for this is also interesting, for it draws on another familiar argument against the possibility of moral objectivity which I mentioned earlier. He points to a significant difference between factual utterances and moral ones, the difference being that moral utterances are action-guiding in a way that factual ones are not. It is not clear that there need be such a difference. Statements such as 'this package is fragile' or 'major road ahead' are certainly meant to be action-guiding; and statements such as 'it is your duty to desist', though perhaps grammatically dissimilar from 'you ought to desist', clearly serve to convey the same meaning. I do not want to deny that there may be interesting syntactic differences between factual and moral judgements, but these differences are beside the point here. They are beside the point because whatever the logical or grammatical status of moral judgements, it is still the case that when a man makes a moral judgement he expresses an action-guiding belief or conviction which he has. Our interest is not in the logical or grammatical status of what has been said, but in whether his belief or conviction has any justification beyond the fact that <u>he</u> happens to have that belief rather than some other one. In particular, the problem about objectivity induces us to wonder whether there might be some <u>generally</u> valid justification for his moral belief, such that anyone who failed to believe that <u>that</u> is the right way to act, or that <u>that</u> is morally good in itself, would in some sense be in error. Putting the matter in this way raises no question of anything existing or not existing; and though I have yet to show that objectivity can be presented like this, the subjective-objective antithesis will not after all be out of place in Ethics if it can be. The assumption that objectivity has no place in Ethics which Ayer makes, stems from the assumption that objectivity

must be tied up in some way with the existence of certain kinds of object. From there it is an easy step to suppose that if existence can only be asserted by making statements, any field of discourse which does not lend itself to the making of statements can have nothing to do with objectivity.

In order to criticise the view of objectivity which links it to the issue of the existence of something, and also in order to begin to develop my own view of what it means to call a judgement objective, I shall turn to a more explicit attempt to state what moral objectivity amounts to. The following passage is from Professor Edward Westermarck's book 'Ethical Relativity', written by a man who was first of all an anthropologist, and written during the heyday of scientific approaches to moral philosophy. It is typical of that approach in that it treats the empirical sciences, with their verificatory distinction between truth and falsity, as the model of what constitutes objective human thought, as well as asserting that objectivity involves the existence of something. Thus it combines both the views I have so far been criticising, although in the passage I have chosen to quote there is no specific mention of of facts or verification. Westermarck's opening paragraph runs as follows:

"Ethics is generally looked upon as a 'normative' science, the object of which is to find and formulate moral principles and rules possessing objective validity. The supposed objectivity of moral values, as understood in this treatise, implies that they have a real existence apart from any reference to a human mind, that what is said to be good or bad, right or wrong, cannot be reduced merely to what people think to be good or bad, right or wrong. It makes morality a matter of truth and falsity, and to say that a judgement is true obviously means something different from the statement that it is thought to be true. The objectivity of moral judgements does not presuppose the infallibility of the individual who pronounces such a judgement, nor even the accuracy of a general consensus of opinion; but if a certain course of conduct is objectively right,

it must be thought right by all rational beings who judge truly of the matter and cannot, without error, be judged to be wrong."*

Westermarck goes on to deny that Ethics is a normative science, and from the passage just quoted it is not hard to anticipate why. He denies that Ethics is a science because he denies that there are any "moral principles and rules possessing objective validity"; and this in turn depends on the denial that there are any moral judgements which are known to be true. This involves a straightfoward identification of objectivity with truth, and, what is more, with truth as established by methods appropriate to the establishment of truth in the empirical sciences. Thus Westermarck's notion of objectivity resembles Emmet's in that it really amounts to an equation of objective judgements with factual judgements. But in concentrating on truth as he does, Westermarck rather suggests that for him a judgement will only qualify as objective if it is actually true, whereas Emmet's argument seems to have equated objectivity only with the capacity to be verifiable. Neither writer, however, is very clear about this. In the same vein Westermarck writes later: "If there are no moral truths it cannot be the object of a science of ethics to lay down rules for human conduct, since the aim of all science is the discovery of some truth...If the word 'ethics' is to be used as the name for a science, the object of that science can only be to study the moral consciousness as a fact.**

Scientific Ethics, then, would really be a branch of psychology. And elsewhere Westermarck makes it clear that he thinks that moral judgements are grounded in the emotions, and are therefore expressions of a species of personal feeling; and that consequently there could be no moral judgements which could possibly count as true. (This, he thinks, is because judgements based on personal feelings will not be

* 'Ethical Relativity', Littlefield Adams, 1960 edition, p.3
** op.cit. p.61

generally agreed on, and cannot be the subject of well-founded agreement. In view of my earlier remarks on the impossibility of agreement over expressions of preference, Westermarck is here understating his case. Strictly speaking, what Westermarck thinks a moral judgement is could not be the subject of any sort of agreement.)

Westermarck, then, is contrasting the "supposed objectivity of moral values" which "makes morality a matter of truth and falsity" with "merely...what people think"(i.e. with opinion, a word he uses later). And clearly he is treating these two categories of judgement as mutually exclusive and exhaustive, which, as I have already suggested, is a mistake. I shall deal with the relationship between objectivity and truth at some length in the next chapter. Here I want to consider his other assertion, that the objectivity of moral values would involve their having "a real existence apart from any reference to a human mind".

Now if this assertion is intended to imply that the objectivity of moral values requires there to be an existent, or perhaps subsistent, realm of shadowy moral beings, it is an extremely dubious claim, as Westermarck was, of course, well aware. Yet his use of the word 'real' indicates that he did believe this to be a necessary condition of there being moral objectivity. However, moral judgements could be objective without there being any such realm of existents, a fact which is made more obvious if we accept for the moment Westermarck's idea that to make an objective judgement is to all intents to make a true statement. Now there are a very large number of true statements which can, if we choose to put it that way, be said to exist apart from any reference to a human mind. (There are, for that matter, a very large number of false ones too.) I take it that a statement or a judgement exists with reference to a human mind insofar as someone is, or has been, or will be, entertaining the ideas that the judgement or statement expresses. So no doubt there will be some statements and some judgements which never have been thought of and never will be. To

regard them as existent is to express in a misleading (because categorical) way a hypothetical claim about what would be required of us *if* we were to think of them. *If* we were to entertain a true proposition or make an objective judgement, such a proposition or such a judgement would have to fulfil certain conditions in order to qualify as true or as objective, and that is all. There is no requirement in these conditions that anything should have to exist, although this will be the case if we are considering true statements about the material world. That, however, follows from the meaning of 'material' and not from the meaning of 'true'. So far as objectivity itself is concerned, Westermarck himself specifies what amounts to a set of conditions which a judgement must satisfy if it is to count as objective. A course of conduct, to be objectively right,"must be thought to be right by all rational beings who judge truly of the matter and cannot, without error, be judged to be wrong."(ibid.). I do not myself agree that these are the necessary conditions for a course of conduct being objectively right (not least because I do not think that judgements about what is right *are* objective); but that, for the moment, is beside the point. Here we are interested in whether the conditions for objectivity require anything to exist, and a little reflection on Westermarck's own specifications for moral objectivity shows that the existence of anything is not a requirement.

Someone might object that my characterisation of an objective judgement as a judgement which *would* satisfy certain conditions *if* anyone made the judgement, does not meet Westermarck's requirement that objective moral values would have to have a real existence apart from *any* reference to a human mind. For this analysis still makes reference to a human mind, if only hypothetically. To take this line, however, would be to make an absurd requirement, and one that was surely not what Westermarck intended. For what species of existence involves the denial of the *possibility* of being thought about by someone or other at some time or other? Such existence would not be knowable at all. Westermarck only meant that object-

ivity imposed more stringent conditions than merely the requirement that someone happens to believe something, or even the requirement that everyone does. Any confusion here is brought about because Westermarck has presented what is really a set of conditions for justified assertion as if it was an ontological claim that was being made. As should be quite apparent by now, this is quite easy to do as long as the paradigm for objective judgements is taken to be the factual statement itself, where the satisfaction of truth-conditions does entail the existence of some thing or some state of affairs.

Before moving on from Westermarck's notion of objectivity, it would be as well to mention one important respect in which what he says is correct. He is quite right to recognise that "the objectivity of moral judgements does not presuppose the infallibility of the individual who pronounces such a judgement, nor even the accuracy of a general consensus of opinion", (loc.cit., my stress). He recognises that to call a judgement objective is to say more than that there is merely a coincidence of individual opinions on the subject in question. Now clearly we do not believe that what is right (as opposed to what is thought to be right) is determined by majority decision. It is not self-contradictory to assert that in moral matters, as in others, the majority is wrong — though, as Professor Hamlyn has pointed out, it would be odd to regard the majority as always or nearly always wrong without being able to give some further account of why this should be so.* The conditions for the objectivity of a judgement, then, amount to more than the fact that a majority of people would make that judgement or assent to it, or even that everyone would. What these further conditions are will begin to emerge in the next chapter. It is enough to reiterate here that they are not the requirement that anything should be true in the way that factual judgements are true, or that anything should exist.

* Hamlyn, op.cit.

So far I have argued that the issue of objectivity has been confused with questions about factual judgements, truth, and the existence of objects. It has not been altogether easy to keep these different confusions separate, partly because the notion of objectivity is itself still unclear, but partly also because these confusions all stem from the same general source, and exhibit a degree of mutual dependence. They are all products of the assumption that objectivity belongs in some way exclusively to those areas of human experience which the physical sciences can describe; and no doubt the success and prestige of the physical sciences has done something to reinforce this attitude in various ways. But although moral philosophers are very prone to contrast the status of value-judgements with the status of judgements about the material world, where things seem more straightfoward, it is far from obvious that what it means to call an empirical judgement objective is any better understood than what it would mean to call a moral judgement objective. I mentioned earlier that the anti-objectivist argument brought against the Ethical Intuitionist might equally well be brought against the naive realist. For if it is denied that any moral intuition could be objective, because the intuition can only be inspected by the lone individual who claims to have it, we might ask on what basis it is claimed that the ordinary man's personal perception of the contents of the material world entitles him to claim that at least some of the judgements he makes on <u>this</u> basis count as objective. In other words, the basis for claims about the objectivity of empirical judgements is just as much in need of investigation as any claim about the supposed objectivity of moral judgements. And it will not do merely to equate objective judgements with judgements about facts, or judgements about objects, or judgements which are true, as if there were no problems there. For quite apart from any other problems there are about our knowledge of the external world, objectivity is <u>not</u> the same thing as truth or existence or being a fact. We must therefore look more closely than is customary at the relationship between

objectivity and truth, and at the real basis for saying that empirical judgements can be objective. This will occupy much of the next four chapters of this thesis, and only then will we really be in a position to decide whether, and to what extent, moral judgements could, after all, be objective.

The assumption that among non-analytic judgements it is only empirical judgements which could count as objective has had one other important consequence for moral philosophy, which it would be as well to mention here, although it will receive much fuller treatment in its own chapter (Chapter Six). We have seen that it is widely assumed that since objectivity is the preserve of empirical judgements, for one reason or another, moral judgements could not be objective in their own right. Any possibility of there being objective moral judgements, therefore, would depend on there being ways of arguing from (objective) facts to moral conclusions with the full force of logical entailment. Now one important use that has been made of the so-called Is-Ought argument has been to show that moral conclusions cannot be deduced from factual premisses alone, and therefore moral judgements could not be objective. Whether or not this argument is wholly correct has been the subject of lengthy dispute. But whether or not it is correct, the above conclusion only follows on the unspoken assumption, with which we are now thoroughly familiar, that only empirical judgements can be objective in their own right. Therefore, as I shall argue later in detail, the Is-Ought argument is decidedly suspect as a proof of the non-objectivity of moral judgements. The belief that the argument does prove this is just one more example of the ready assumption that to call a judgement objective is to say that it has the features which only empirical judgements really have.

CHAPTER TWO

OBJECTIVITY AND TRUTH

In Chapter One I said that the concept of objectivity has regularly been confused with a number of other notions, and that in consequence the possibility of moral objectivity in particular has been denied for inadequate reasons. The confusions I described there — such as the confusion of the objective with the factual, the existent, the empirically verifiable and the scientifically accessible — possess a good many subtle interconnections and have a corresponding tendency to reinforce one another; and beyond what I have said in the previous chapter, I do not intend to separate them or to try to deal with them individually. For they will all collapse together when the real basis for empirical objectivity eventually becomes clear. But we shall not see right through the illusion they create of being the bases, or even the defining cases, of objectivity until we reach the mid-point of this thesis, when we will have explored more fully the reasons there are for saying that judgements about the material world and its contents are objective judgements.

However, there is one feature of most, if not quite all, of these confusions about the real basis of objectivity which has a particularly strong influence on our thinking, and might be regarded as central to the confusion I have in mind. This is the virtual identification of objectivity with truth. We have already encountered the verificationist approach to objectivity, which regards a judgement as objective only if it is capable of being true, and if its truth can be established by verificationist methods. And we have encountered Westermarck's view that if a judgement is objective, questions about its objectivity are, in effect, questions about its truth or falsity. As we have seen, the conflation of objectivity with truth leads to various forms of knock-down argument against the possibility of <u>moral</u> objectivity, just because moral judgements are not supposed to count as either true or false. It is hardly surprising that many philosophers

should have found it so easy to swallow this line of argument. For, insofar as we have one at all, our vague and shadowy paradigm of an objective judgement is one of a judgement about the external world which is, as it happens, invariably also true. When, for example, I look at a cow and judge that it is bigger than a rabbit, the objectivity of my judgement seems to be confirmed by the judgement's obvious truth. It is an easy step from noticing this apparently constant conjunction to calling the judgement objective <u>because</u> it is true, and from there it is easy also to take the further step of equating truth and objectivity completely. The main aim of this chapter is to argue against the taking of those easy steps. Objectivity and truth are not equivalent to one another; nor is the objectivity of a judgement dependent on its truth. For a judgement can be objective without being true, and true without being objective. Objectivity and truth are indeed closely connected concepts. Yet for all that, they are distinct concepts; they are not co-extensive and they have different functions to fulfil in the structure of our thought.

Let us begin by considering some of the requirements a judgement must meet if it is to be objective. It will be recalled that in the previous chapter I argued against the belief that moral judgements could not be objective since they only expressed an individual's personal view, and so were akin to judgements of taste or expressions of personal preference. This view of moral judgements was mistaken, I said, because, like some other kinds of non-factual judgement, moral judgements display features incompatible with their being wholly dependent on the individual's own beliefs or attitudes. However, the argument possesses some residual interest for us because it draws attention to one basic requirement for objectivity. If a judgement is to be objective, it cannot be of the sort where grounds for holding it are entirely dependent on the individual, or indeed one where it is improper to speak of there being <u>grounds</u> at all, and misleading to call these things <u>judgements</u>. (This is the case with avowals, for instance. I shall continue to

speak of judgements in such cases only for convenience, and where it does not obscure what I mean to do so.) We can properly speak of objective or subjective judgements only where there is scope for real agreement or disagreement between individuals; and that can occur only where it is possible for different individuals to make the same judgements as one another, and not just their own personal, self-justifying 'judgements' about the same topic. So a minimal condition for even considering a judgement to be objective is that it should be possible for people to agree about it.

Often, where people can agree in their judgements, they will agree. However, as we noticed in connection with Westermarck's views in the previous chapter, the fact that others share my view is not in itself a reason for calling my view an objective one. For if the majority, or indeed everyone, happen to judge as I do about some matter, their agreement does not suffice for objectivity. It is possible to say, at any rate in particular cases, that everyone's view is warped or biased or unfounded. And in any case, the fact that we all agree in making a particular judgement is no guarantee that we are all making it for the same good reasons, or for any reason at all, good or bad. Some reference to the grounds for making a judgement is inescapable when its objectivity is in question. What would we say about a blind man who, when placed before an unfamiliar object, was unfailingly able to make the same judgement we would all make about its colour? We would not be very willing to call his judgements objective, just because the blind man lacks the faculty normally employed in making such judgements, and therefore cannot be said to have the proper reason for making the judgements he does make. But suppose the blind man later admits that he always guesses in these trials? There is then no basis at all for calling his judgements objective, even if they are right, since he admits that he has no reason at all for saying one thing rather than another in those cases. I am not suggesting that there can never be a reason for guessing one thing rather than another. But to call one's reason a reason for guessing is to imply that it is not a reason for knowing or a good reason for believing.

It seems evident, then, that a major requirement for the objectivity of a judgement is that it should be made with some sort of justification, or that it must reach some level of adequacy, or that it must meet some standard of competence that guarantees that it has been made with good reason. These requirements are here expressed in rather vague terms, as indeed they were in Chapter One, where the idea first appeared. This vagueness stems partly from the fact that in some respects the requirement will differ according to the type of judgement that is being made, and partly because I do not want to deal with the absolutely general basis of justification in judgements until a later and more appropriate stage of my argument (Chapter Five). It has been enough up to this point to speak of reasons and grounds, good reasons and good grounds, and justification, sufficient grounds and criteria, without much differentiation except to indicate (a) the presence of <u>some</u> degree of justification for the judgement in question, and (b) roughly the strength of the justification involved. Now I will try to be a little more exact.

It is to be expected that not just <u>any</u> degree, or just <u>any</u> type of justification will be enough to render a particular judgement objective, and that this is not a matter which the individual can decide for himself. The first and obvious condition to be met if a judgement is to be <u>sufficiently</u> justified to count as objective is that an appropriate method of judgement should be employed. This involves two things. One is that an appropriate <u>faculty</u> of judgement should be employed; that is, the appropriate sense or combination of senses should be used in making the judgement. The other is that due regard to the <u>circumstances</u> in which the judgement is made is required, because there are certain circumstances which <u>demand</u> to be taken into account if a judgement is to be properly made. By way of illustration, let us turn to the cow and the rabbit once again. A judgement about the colours of these two creatures <u>must</u> involve the use of one's eyes, since vision alone is the determinate faculty for judgements

of colour; whereas a determination of relative size may involve either or both of two faculties, vision and touch. But if we judge relative size by visual means, relative distance of the objects from our eyes must also be taken account of, as a circumstance necessary to proper judgement, since it is a condition of judgements about size being objective that we recognise that distance changes apparent size but not actual size.

There are a great many such considerations which apply to the various types of judgement that we may make objectively. It would be pointless to try to list all of them, and not just because there are so many, but also because it would be a mistake to think that the set of such conditions is either complete or incapable of any modification. The reasons for this will appear in the course of the chapter, as the nature of the interaction between objectivity and truth becomes plainer. Part of Chapter Three also represents an informal attempt to outline in some detail the conditions for one kind of objective judgement, and there I shall also discuss the possibility of change in the faculty of judgement deemed to be appropriate. So I will not go into more detail here. However, in general terms we may say that a judgement is sufficiently justified, and therefore objective, if it is made by the use of an appropriate faculty and with due regard for the conditions relating to the type of judgement involved, something which depends in turn on the property or properties of the object which the judgement attempts to establish in the examples I have given. These two factors, taken together, constitute a <u>method</u> of judging. And the critical feature of an objective judgement is just that it should be made <u>by a method appropriate to the sort of judgement it is</u>.

The second condition for a judgement's counting as sufficiently justified is that the method referred to above should be such as will result in uniformity of judgement. Thus, if a person has applied an appropriate method of judgement to a situation, and applied it correctly, there will only be one judgement that can be made as a result. For example, the conditions for judging the relative sizes of a cow and a rabbit must be strict enough so that

I will in the same circumstances always make the same judgement if I judge according to the method. I will always judge that the cow is bigger than the rabbit, or else always judge that the rabbit is bigger than the cow, or else always judge that the two are the same size. The method of judgement will not be adequate to produce sufficiently justified judgements if it sometimes allows me to reach one conclusion when I follow it, and sometimes another. The corollary of this uniqueness of resulting judgement is that the judgement should be universally made in that way. In other words, it does not matter whether it is you or I who makes the judgement. In a particular set of circumstances, if we have paid due attention to the conditions of proper judgement and employed the appropriate faculty, the results of the judgement should be the same, no matter which of us makes it. None of this will guarantee that any particular judgement you or I may make will actually be true; nor does it mean, of course, that we will always make our judgements properly, and so agree in the conclusions we reach in every case. Where we do not, there is an indication that either one of us at least has deviated from the appropriate method of judgement, or else that the method itself is inadequate as a guide to the truth in this particular case. In either case, we will look for an explanation, because there must be something unusual in the situation, or something unusual about the approach to it that one of us has taken. Where this does not happen, as often it will not, our method of judging has provided us with a reliable guide to the truth of the matter. But the truth in any particular case should not be confused with the method by which we have come to it. A judgement may be sufficently justified, and therefore objective, even if it fails to reach the truth, provided the judgement is made in accordance with a method which satisfies the conditions laid down above.

Although all this might seem obvious enough when we are making judgements about the material world and its contents, we might wonder whether there is always just one justifiable way to see things of other sorts.

Could there not be areas of human thought and experience where a number of different ways of seeing some thing or some situation are equally acceptable and appropriate? In a way this is correct, since we sometimes do recognise that someone's opinion is justified up to a point. But strictly we should see this as only justification from a particular point of view. To speak of justification simpliciter, and even more to speak of sufficient justification, is to imply that there is a single standard of judgement which applies to the cases in question. We cannot say that two incompatible views are both equally correct, in any context where the word 'correct' has a real application. And we cannot call two differing views both sufficiently justified just because there is something to be said in favour of each of them. No field of judgement, then, can contain genuinely objective judgements unless it is the case that if everyone judged properly in these cases, they would all judge in the same way, with the same result.

It is through a mistaken appreciation of this last point that the concept of objectivity has been so often and misleadingly linked to the concept of necessity. Some philosophers have given the impression that if a judgement is objective, it must be logically necessary for everyone to reach the same conclusion in judging.* This seems a little harsh on empirical judgements, which are traditionally thought of as the main or only field where objectivity occurs, but which have rarely been thought of as necessarily true! Other philosophers have sometimes implied that a sort of natural necessity, a feature of our nature, drives us all to see things in a certain kind of way, and so accounts for objective judgement.** This is less drastic, but no less misleading. For it makes our tendency to see things objectively a matter of fact about us, and would mean that a person who failed to see things

* See Appendix

** See Chapter Four

as the rest of us do would be regarded as odd, no doubt, or perhaps unnatural, but could not be regarded as wrong in his judgements, unless we crudely treat 'wrong' as equivalent to 'different'. However, if we need to speak of necessity at all in connection with objective judgements, the only necessity we would need to recognise is what we might call 'rational necessity', the requirement that in making, or trying to make, an objective judgement, we are committed to judging according to rules which establish which judgements count as objective and which do not. These rules, of course, are the conditions for sufficient justification of which I have already spoken. Westermarck expressed something of this when, in the passage quoted in Chapter One, he wrote: "If a certain course of conduct is objectively right, it must be thought right by all rational beings...and cannot, without error, be judged to be wrong."* No-one, then, is compelled by logic or nature to be objective. If they are not, the penalty is error, not self-contradiction or unnaturalness.

Now we can see the features of objective judgement which I have been discussing exemplified in practice in just those cases where the individual's claim to have been objective is called into question. When two people cannot agree on whether some object before them has a particular physical property or not, they will often turn to a third party in order to have their particular judgements confirmed or refuted. Two requirements are made of the third party. The first is that he will make an honest, unbiased judgement. (Setting aside questions of personality, this is an easy requirement in nearly all cases of empirical judgement, but becomes important whenever a man might have an interest in one answer rather than another being given.) The second, and as a rule the more important consideration, is that the arbiter should be at least as capable of making the judgement objectively as the disputants. This means that he is not disbarred by lack of the appropriate faculty or by lack of understanding of the circumstances which must be taken into account if the judgement is to be made properly from judging aright. The point to notice here,

* Westermarck, op.cit. p.3

though, is that within the boundaries imposed by these demands, absolutely anyone will serve to be an arbiter to such a dispute. This brings home the qualities of uniqueness and universality which we now expect to be a feature of objective judgements. We cannot always tell from what our impartial judge says that he has judged as he has for the proper reasons. But in cases which only involve sensory inspection this is covered by his not being disbarred from the ordinary use of his senses, and we can discover whether he has taken account of what we think are the relevant conditions for making the judgement properly by asking him. We would not have asked a blind man to arbitrate in a dispute over the colour of something, nor a stupid one to arbitrate in a judgement of great complexity or one which required special knowledge. But providing our arbiter is competent to judge, and providing the circumstances are such that judgement is really possible, we can reasonably expect anyone who manages to avoid error to be in a position to settle our dispute, and give the one answer which is actually the correct one.

This being so, we might wish to call anyone not peculiarly disbarred from making judgements of some sort a <u>standard judge</u> with respect to that sort of judgement. It is probable that we are nearly all standard judges with respect to most types of judgement most of the time, at least for non-specialist judgements about the physical world. However, it might be thought that there is an element of circularity in this whole approach to objectivity; for might we not define standard judges as those who judge as the majority judge, and regard anyone who judges differently as a person showing evidence of being disbarred from proper judgement? It might seem that all this is a way of re-introducing the notion that objectivity is to be equated with the majority view. But this suspicion disappears when we also realise that we all lapse from time to time from the standards required for objective judgement, and that it is possible for anyone, inclusive of those forming the majority, to be wrong from time to time. If I recognise that the majority might be wrong <u>in a case where I also might be wrong</u>, I cannot be

appealing to the disparity between my view and that of the majority as a reason for my belief in my own fallibility. In these circumstances, the possibility which I recognise of my being in error must involve comparison of my judgement to a standard which is independent of the beliefs of any individual, or group of individuals, even if on occasion the group includes us all. This belief is not incompatible with my believing also that most people's judgements meet the standard needed for objectivity in most cases. So the appeal to a single universal standard of adequacy in judgements is not an appeal simply to the majority view. It is an appeal to a standard that is independent but generally attainable, which is not the same thing at all.

Perhaps it seems that all I have so far said about conditions for the objectivity of judgements could also be said about conditions for their truth. However, this is far from being the case. A true judgement need not be objective, and an objective judgement need not be true. We can see the former quite easily by reminding ourselves that a man may judge truly of some matter on quite inappropriate grounds, or on no grounds at all. In such circumstances we would not say that his judgement was objective, even if the conclusion he reached was the one he would have reached if he had been objective in his judgement. Even if our conception of truth were to be taken as identical with the judgement arrived at by anyone judging objectively (which it is not, as later arguments will show) it would still not follow that if a particular judgement were true, it was ipso facto objective. For as we have seen, a judgement only counts as objective when it is made with sufficient justification. And a man may guess the truth with little or no justification for his guess. Here I think it is quite obvious that, in cases where we can speak of the truth, to speak of being objective is to say something about the standards for the attainment of truth, rather than about the truth itself.

It is rather more complicated to show that a judgement could be objective without being true, and this is best

done by way of examples. Let us first of all imagine the case of a man who finds a rather nice, but dusty, landscape painting in his deceased great aunt's attic, and wonders from the look of it whether it might be valuable. It looks rather like a Constable, he thinks. The painting is unsigned, but let us suppose that this is no bar to its being authentic. Let us also suppose that there is no previous record of the painting's existence, and no known evidence that Constable ever painted such a picture, but equally there is no clear reason for saying that he couldn't have painted it. So evidence for the painting's probity will have to be internal evidence from the picture itself. The owner no doubt will take the painting to an art dealer, who will have recourse to the art experts. In our imaginary case, the experts are unanimous in their view that the painting displays every sign characteristic of the best of Constable's work, is not one of his son's imitations, is certainly of the right age, shows no sign of being even a contemporary forgery, and so on and so forth. In short, the internal evidence for its being by Constable is overwhelming.

Now here the experts are being what I would want to call objective in their judgements. That is, they are judging in a competent manner, with due regard for the appropriate considerations affecting this sort of judgement, and arriving at the conclusion that anyone with sufficient knowledge and expertise would have to arrive at too. In short, they have every reason for judging as they do. And if they have every reason, and if also there is enough evidence to warrant a judgement without reservations, they have sufficient justification for their judgement.

(I make the experts unanimous in this case for simplicity's sake. But lest it should be thought that unanimity is a precondition of objectivity in such a case, let me emphasise that if the experts were to disagree, as experts often do, that would show one of two things. Either there was, as it happened, insufficient evidence to warrant us in calling their judgements objective, because on such slender evidence no-one is in

a proper position to judge in the full sense I have given to the notion of objective judgement; or else the disagreement would show that some at least of the experts on this occasion were being less than competent, and had failed to be objective in their judgements. Since this is a difficult case, the experts may all have been *trying* to be objective, but failed to be. One of the features associated with expertise is the capacity to recognise how far one can judge on the basis of available evidence, and to go no farther than that, while at the same time having a good capacity to extrapolate what is likely to be the case from the evidence available. In practice, disagreement among experts in this sort of case is likely to be about the degree of assurance that they will give to the judgement that the painting was by Constable. *Ex hypothesi*, however, neither of these factors plays a part in my example, because there is no disagreement among the experts; but the fact that disagreements do occur in real life does not vitiate the point behind the example.)

Yet are we to say, on the basis of the objective judgement that the painting was by Constable, that the statement 'this painting is by Constable' *must be true*? Surely we would not talk of truth in a case like this unless there was a rather different source of evidence available, such as documentary proof that the painting was by Constable. But equally, in the absence of such proof, we do not want to deny that the experts' judgements were objective. For if proof in some form were suddenly to appear, that would not make the experts' judgements any more objective than they had previously been. Judgements do not *become* objective when they are confirmed; they are perhaps *shown* to have been objective to anyone who doubted it before. But the judgements *were* objective all along, by virtue of the way in which they were made.

The objection might be made: but does not the fact that the statement 'this painting is by Constable' *is true*, play a crucial role in the objectivity of the experts' judgements? Isn't it that their judgements

are objective only so long as the painting really is a Constable, even if no-one knows it at the time? But this objection misses the point. The point is not whether the judgement is in fact true (although in cases where it is appropriate to talk of truth we can expect objective judgements to turn out to be true); the point is the matter of how the judgement has been made. Let us stretch our example further and suppose that after the experts have pronounced the painting to be a Constable, incontrovertible evidence comes to light which proves that the painting is really a masterly forgery, one so brilliantly executed that no expert could have detected it from the painting itself, for the forgery was so good that the painting is literally indistinguishable from a genuine Constable. So the experts' judgements are wrong after all. Does it follow that they were not objective? On the contrary. Their judgements would only lack objectivity if they were made in a way which would not normally be sufficient to let them make the justified claim that the painting was genuine. In other words, if no competent judge could have avoided being misled in such a case as this, the judgement he makes may well be objective even though it is false. The idea that someone might be wrong in his judgement, but that, if he made it in the right way, with a proper concern for reaching the truth, it is nevertheless an objective judgement, is central to understanding the difference between truth and objectivity. For if, as I have suggested, objectivity is a function of the making of judgements according to a standard of adequacy which counts as sufficient justification for so judging, then there is no reason why we should not speak of judgements being objective in areas where we should hesitate to speak of judgements as being true or false — providing, of course, the standard of adequacy does apply to judgements of that kind.

Now if cases like the one above can occur, showing that a judgement can sometimes be objective though false, we might expect to find that this will not be an

occasional occurrence but a regular feature of some types of judgement where normal methods of judgement regularly fail to lead to the truth. Such judgements will form a small but significant sub-class of a larger class of judgements where objective methods of judgement <u>are</u> a guide to the truth. Among visual judgements, i.e. judgements made entirely by looking and interpreting what is seen, there is a significant group of objective but false judgements, and this characteristic leads us to call the subjects of those judgements <u>illusions</u>. (I use the term as Austin used it, to refer to cases where what we are inclined to say we see amounts to a mis-description of what is really there, but without implying that we must be fooled by what we encounter into believing that it is really something that it is not.)* The point about illusions is that they are not the result of lapses or failures on the part of the individual. They occur through the failure of a particular method of judgement to lead to the truth in special cases, while still being a reliable guide to the truth as a rule.

The Müller-Lyer Illusion provides a convenient example of this.

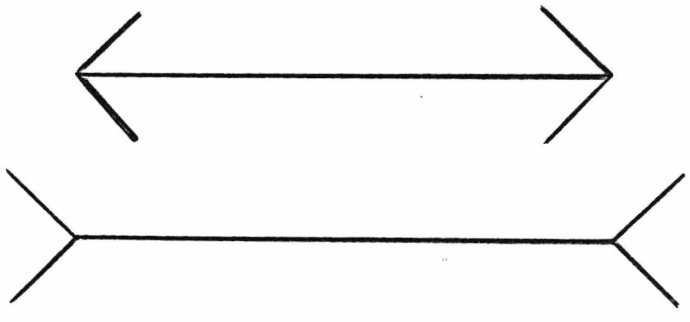

MÜLLER-LYER ILLUSION

* See J.L. Austin, 'Sense and Sensibilia', OUP 1962, sect. 3

When we look at the diagram without being put on our guard by knowing in advance that there is something unusual about it, our visual inspection leads us to judge that one line is longer than the other, whereas the two lines are really the same length. Notice that we are naturally inclined to say that one line <u>is</u> longer than the other, and not just that one line <u>looks</u> longer than the other; for though the latter statement is true, we should only say this if we had some reason to doubt that how the diagram looks is a good indication of how things really are. But we always need a special reason for having doubts of this sort; and until we learn to recognise cases like this as being in some way exceptional, we shall go on saying that the lines <u>are</u> unequal in length on account of the way they look. The point to notice is that there is nothing unusual, and nothing subjective, in such a response to seeing the diagram. If, distrusting our own eyes or our own judgement, we were to enlist the aid of standard judges, we should find that their judgement coincided with ours — unless, of course, <u>they</u> have some doubt about the normality of the example. There is no degree of heightened visual competence which will result in seeing the diagram differently. In short, we see the diagram as it is normal to see it, and we judge as it would be normal to judge when we say that the lines are different in length. Indeed, if this were not the case, and if moreover there were not a strong tendency to see the lines as unequal even when we know that this is an illusion specially designed to produce a misleading judgement on our part, the illusion would not be psychologically informative at all, and would find no place in psychology textbooks. So here I want to say that the judgement that the lines are different in length, though false, is an objective judgement nevertheless.

The reason why the judgement is objective is quite straightfoward. How things look is normally a sufficient basis for saying how things really are in such cases. We must see the failure of our judgement to be true in this particular case against the background of our regular success in basically similar cases.

But surely, it will be objected, the judgement in question here cannot be objective, even if it is one that people will all make unless they realise they are confronting an illusion. For this judgement contradicts another judgement, the judgement that the lines only *appear* to be unequal but are really the same length. And this judgement is not only itself objective, but also true. Its objectivity is established by the very processes which lead us to recognise the visual illusion for what it is.

In fact, though, the two judgements only contradict one another in respect of their truth-claims about how things really are. That does not mean that they cannot both be objective. For objectivity, as I have argued, is concerned with the manner in which a judgement is made, and not with its truth or falsity *simpliciter*. If the unsophisticated judgement that the lines are unequal in length was properly made, it is objective. The more sophisticated judgement does not vitiate this objectivity, but modifies the value of the unsophisticated judgement as an expression of the truth, showing that in some cases at least there is a limitation on the utility of simple visual inspection as a means to discovering how things are, and confining its capacity to express the truth, in this case, to the truth about how things *appear* to be. We would of course have to discover the need for this limitation on visual inspection as a means to the truth by some other objective method of judgement which we believe to be more appropriate (because more accurate). This is a perfectly reasonable procedure, however; for the point of my argument here could be put by saying that an objective method of judgement is not an infallible method.

A further objection might be raised, in the following terms: surely the claim that 'the lines are unequal in length' is objective, must be wrong. For surely it is ruled out by the condition imposed earlier in this chapter, that a judgement is only objective if it is made with sufficient justification. But how things look is *not* in this case a sufficient justification for saying

how things are, since how the lines look is not how they are. The only cases where such justification could be sufficient would be cases where things actually are the way they look to be. And this goes to show that objective judgements, as I have described them, really are dependent on truth, since only in cases where the judgement is true will the justification turn out to have been sufficient. However, we can see a way past this objection if we turn for a moment from the logical features of 'sufficient' to the epistemological features of 'justification'.

As we have already seen, we are only entitled to call a judgement a justified one if it is made with good reason according to some acceptable method of judging. And in a sense, as we have seen, the question of the truth of the judgement is irrelevant here, since a wild guess may be true but is by definition unjustified as a judgement. Now a method of judging cannot be a <u>method</u> if it can apply only to a single case of judgement (although no doubt there are instances of a single judgement becoming a paradigm case for some reason, and so establishing a method of judgement for similar cases). I judge, for instance, that the paper on which these words are written is white. That judgement is grounded on an inference from its present appearance to me. But it could not be so grounded if every fresh appearance of a sensation of whiteness required a fresh sort of justification for my being able to judge that here <u>is</u> something white. This being so, we can always expect there to be cases where the justification will not ensure true judgement, since the justification will always be a general one, but the question of whether the truth-conditions for the judgement are fulfilled will be particular to each separate judgement. Furthermore, it will make no difference if we start to write qualifications into our method of judgement, saying, for example, that the method only holds good in certain lighting conditions, and so forth. To an extent we have to do this anyway, insofar as each of the general types of judgement I spoke of earlier needs us to take account of the conditions for judgement so as to be sure they are such that the judgement may be made properly.

Obviously, the set of conditions for objective judgement is capable of refinement and sophistication. But the method will cease to be a method (as well as becoming too unwieldy to be usable) if we write into it so many qualifications that it finally becomes the disjunction of the particular truth-conditions of all the judgements to which it could ever be applied. (And there will be infinitely many such judgements.) Thus the notion of a sufficient justification could never be one which would result in every judgement to which it applied being true. It could only be the sort of justification which was more capable of resulting in true judgements (where it is appropriate to speak of truth) than any other sort of justification which a standard judge has available. We shall see in a later chapter that this is rather a clumsy way of putting the matter, but I think it will do to meet the present objection. The notion of sufficient justification was introduced to express the idea of a unique and universal standard of correct ways of judging. It was not a requirement of that notion that such judgements should always be true. It is enough if these methods of judgement are appropriate as ways of getting at the truth. It will be apparent that such methods cannot as a rule lead us astray and still be appropriate methods for reaching the truth. And this is another reason why objective judgement is so readily confused with true judgements, and especially with true empirical judgements, where the procedures appropriate for reaching the truth are so readily understood and applied, and are so regularly successful. Objective judgements will normally turn out to be true ones, and judging with sufficient justification will normally result in our reaching the truth. But the fact that we do not always succeed is no bar to our having faith in our methods, providing our methods are normally adequate.

So far in this chapter, my main concern has been to indicate that objectivity and truth are distinct, as a preliminary to showing that the real basis for the objectivity of judgements lies elsewhere than in their truth, and in order to free objectivity from being too closely

associated with true empirical judgements. Thus I have argued that to call a judgement objective is to say that it conforms to a unique and universal standard of proper judgement, which will normally lead to the truth where truth is attainable; but that this is by no means the same as saying that objectivity is equivalent to truth, or that the objectivity of a judgement follows from that judgement's truth. The matter might be summarised by saying that objectivity is the general guarantee of a judgement's fitness to express the way things are, whereas the claim that a judgement is true is the indication that a judgement has succeeded in doing this in the case to which it actually relates.

Having made this distinction, however, I cannot leave so important a topic without saying something more about the relationship between the two concepts I have been at pains to separate. This is all the more important since my eventual concern will be with the question of the objectivity of moral judgements, where the applicability of the terms 'true' and 'false' is felt to be problematic. Later in this thesis, when our attention is focussed on moral judgements, it will be possible to indicate those types of moral judgement which might without impropriety be described as true. Here, though, the discussion will be more general.

The gap which I have been exposing between objective judgements and true ones helps to point up just how demanding is the conception of truth that we have. For a judgement to be true, it is not enough that it should be believed, even by everybody. It is not enough even that we should have good reasons for believing what we do. For no good reason will ever amount to a logically sufficient reason, at least in connection with synthetic judgements, since the truth of a synthetic judgement could always be otherwise than it is, and therefore always other than the evidence suggests that it is, no matter how strongly suggestive the evidence may be. To be true, a judgement must actually succeed in expressing the way things really are. Now it might be wondered whether such a conception of truth is even feasible, since it seems to

require us to transcend the limits of our sensibility. Yet I think that there is this requirement in our conception of truth, and properly so. For truth represents an absolute or final standard against which our judgements can be assessed. Truth is, so to speak, a limiting concept. It is evident that without some conception of truth there would be no reason for treating an objective judgement as outweighing a subjective one, or believing that a judgement made without justification is in general less satisfactory than a judgement made with justification, since it is as a general means to reaching the truth that reasons for judging become warranted as reasons. Yet without a conception of truth which transcends the possibility that our judgements will give complete expression to its content, we should be left with a picture of objective judgement which I have already shown to be unsatisfactory. For if the truth were finally attainable by means of objective judgement, then a judgement made with sufficient justification would in every case have to result in something which we not only believed to be true, but which we would also have to regard as incorrigible. Our notions of truth and of objectivity would then be reduced to a notion of something identical in its outcome but also something which could not be altered or amended in the light of further observations or further considerations of other sorts. This, however, is far from being the case.

This static model of objectivity and truth, which would make it impossible to distinguish between the truth itself and what we believe to be true for good reasons, has affinities with outmoded inductivist methodologies of science in which human knowledge was seen as a body of confirmed truths which grows by gradual accretion. It also corresponds with more recent conceptions of science which take from Logic the model of an axiomatised and fully deductive system of thought in which the truth would find its final expression, although the finality of such truth would only be finality-within-the-system. Both these views of the nature of scientific truth have been attacked by Popper*, whose theory of scientific

* K.Popper, 'The Logic of Scientific Discovery' in 'Conjectures and Refutations', Routledge, 1963

method emphasises falsifiability rather than verifiability as the criterion of being a genuine scientific hypothesis. Kuhn*, approaching the issue from the standpoint of the History of Science, has also argued against the notion of final truth being attainable by science, although he suggests a sense in which science can be said to progress because the process of replacing theories with other ones is, he argues, irreversible.** I am not entirely convinced by Kuhn's argument at this point; however, we do not need to decide whether, or in what sense, science can be said to approach any sort of ultimate truth. I have only mentioned these attacks on theories of science which represent science as capable of expressing truth in some absolute form because the attacks lend some support to my contention that we must see our beliefs about what is true, however objective our judgements, as corrigible in the last resort. And this corrigibility can be represented as a process of interaction between the methods of judging which we can regard as justified and the well-founded beliefs that we actually hold; that is, as a process of mutual reinforcement and correction between objectivity and the truth-as-we-believe-it-to-be. The alternative, which is to conflate objectivity and truth, belies the conception of truth that we actually have, as well as showing a misunderstanding of the concept of objectivity.

A simple example of the interaction between truth and objectivity is available if we return to consideration of our judgements about the Müller-Lyer diagram. It will be remembered that our (unsophisticated) judgement was that the lines are really of different lengths; but I argued that although this judgement is false, yet because that way of seeing the diagram is the normal way to see it (whether or not we recognise the illusion for what it is), our judgement that the lines are unequal in length (because they look unequal) is objective. If objectivity always led to the truth, we should be forced to conclude here

* T.S.Kuhn 'The Structure of Scientific Revolutions' (2nd edition 1970) Chicago.

** op.cit. especially pp.205-207

that the lines really were unequal in length, even though this judgement can be modified by other, equally proper and objective, but more precise, methods of determination, such as measuring the lines with the aid of a ruler. So, as we have already noted, the fact that we judge objectively does not here guarantee that our judgement has enabled us to reach the truth. Indeed, in this case, the idea that objectivity invariably leads to the truth would be absurd, once we notice that in correcting one judgement by another we are employing what is in a crucial respect the same method of judgement. For in order to discover that the lines really are the same length, we are required to note that the lines on the diagram are both exactly the same length as the distance from the end of the ruler to one particular point on its length. And we cannot do that without employing the principle that looking the same length is sufficient grounds for really being the same length, the very principle that this example throws open to doubt.

In discovering this, we continue to accept the general method of judgement of length by eye, not only because of its utility as a means to the truth in other cases, but also because here it serves as the means to correct the conclusion reached by the original observation of the diagram, and so indicates a limitation on the applicability of the general method of judging by this sort of appearance. Yet we must also employ the method in order to recognise its limitations. The method which here resulted in the error is also the method by which the error is rectified and our conception of what is true enhanced. Here, then, an objective method of judgement first fails to reach the truth, then increases our grasp of the truth; and that in turn modifies our view of the utility of objective methods of judgement in certain cases. None of this would be possible, however, if objectivity and truth were identical.

A more general, and richer, example of this interaction is provided by examples like the belief, once widely held, that the Earth is flat. Although we regard this belief as a false one, and are right to do so, we

cannot merely dismiss it as a vulgar superstition or a silly error. There is, after all, a good deal of straightfoward evidence to support the hypothesis that the Earth is flat and to contradict the suggestion that its surface is curved. To start with, the surface of the Earth looks flat, and we need reasons to suppose that how it looks is <u>not</u> a good guide to how it really is. True, ships seem to disappear below the horizon in a way that makes perfect sense if we believe the surface of the Earth to be curved, and Copernicus for one thought this to be a highly significant example in the context of his theory.* But since water finds its own level, and the sea is not in constant retreat over the horizon, it seems just as sensible to postulate that how things look fails to be a guide to the truth in cases of distant vision (as with ships on the horizon) rather than in cases where we can observe with accuracy (as with experiments to show that water finds its own level).

We no longer believe that the Earth is flat. This shows that a quantity of straightfoward evidence, objectively arrived at, does not suffice to establish the truth once and for all. Yet at the same time, our idea that the look of our surroundings is in general a good and indeed normally sufficient guide to their real properties has not changed. Rather, what was formerly regarded as an aberrant type of judgement, the well-attested observation of ships seeming to sink over the horizon, has had its inherent objectivity vindicated by the explanatory power of a new hypothesis. Thus the objectivity of appearance-judgements as sufficient grounds for saying how things are has undergone change and enhancement, becoming limited in one case but being confirmed in another; but it has not in any sense been generally weakened. On the other hand, our ideas about the Earth's true shape have been radically altered.

* For an interesting discussion of whether or not this is a crucial test of Copernicus' theory, see I.M.Copi's 'Introduction to Logic' (4th edition) Macmillan pp. 449-452.

The foregoing examples might suggest that the types of objective judgement we can make are permanently fixed, and that only modifications to the range of their application can occur. It might not seem as if an entire method of judgement, once regarded as objective, might become suspect. Certainly, we would not expect there to be more than minor modifications to the range of objective judgements we can make based on observation alone, unless there were to be some radical change in our senses, and hence in the kind of creatures we are. (I shall have more to say about this in the next chapter.) However, the more closely a method of objective judgement involves belief in a particular causal hypothesis, the more vulnerable it is to the loss of its claim to objectivity if the hypothesis loses credibility. In a culture possessing very different ways of looking at the world from our own, the notion of a cause is sometimes closely allied to that of a magical influence in some cases which are felt to be in some way special. For example, the idea that casting a spell on somebody causes him harm may be commonly held. Thus the judgement 'he died because a spell was cast on him' might count as objective within a certain culture, not only because it is widely believed, but because an elaborate and elegant structure of reasons may exist for the belief and, when applied consistently, may amount to a theory with a high degree of predictive accuracy and explanatory power. We should regard judgements made on the basis of this theory as sufficiently justified when, from within that culture, the reasons for those judgements are recognised as unique and universal, so that they indicate not only how people within that culture do in fact judge, but how they would regard it as proper to judge. This would require judgements like 'when a spell is cast on X, X falls sick' to normally be true; and when such a judgement fails to be true, there would have to be reasons available within the terms of the theory to account for this. Provided such a theory could be clearly enough stated, we might imagine it standing up well to the tests a Western scientist would devise to test its validity. And there seems to be a certain amount of

anthropological evidence for the existence of such causal beliefs among primitive peoples, beliefs which for them have the status of objective principles.* Here I am not concerned with what, if any, causal explanation of such occurrences can be given in terms that we would find acceptable. It will be enough that the judgements of the primitive people are objective according to the basic criteria of objectivity established in this chapter, and that their account of what actually happens and why comes as close to the truth of things as they can manage. The significant point is that under the influence of Western scientific culture, such a view tends to decline, and as it does so it loses its claim to objectivity. Belief in the spiritual or magical nature of causation comes to be seen as having less predictive and explanatory power in general than belief in physical causation. The consequence is a radical shift in the primitive culture as to what counts as an objective judgement about why things happen as they do, and indeed about what is happening, because the reasons underlying the judgement are no longer seen as sufficient in every, or almost every, case. This account is no doubt over-simplified; but it serves to illustrate the possibility that a method of judgement can cease to qualify as objective in the course of cross-cultural change. The same might be true as a result of historical change too, though changes are liable to be much less marked. But if this argument is correct, it suggests that our beliefs in what methods of judgement are objective need not be any more rigid than our beliefs about what is true.

The foregoing arguments are intended to illustrate the mutability of what we think is the truth at any time, and suggest that what counts as an objective method of judgement both affects and is affected by it. A number of serious philosophical issues about the nature of explanation are involved at this point, but they are not

* for vivid accounts of such beliefs, see M.Mead 'Sex and Temperament in Three Primitive Societies' (Routledge 1977) and Aubrey Menen 'The Prevalence of Witches' (Chatto & Windus 1970)

directly relevant to my argument, which aimed to show only that in any particular case the question of whether something is true and the question of whether a judgement made about it is objective are two separate issues, even though in many cases the answer given to one question will have a bearing on the other. One consideration which I cannot overlook, however, is the possibility that the view of truth which I have been putting forward in this chapter makes it impossible for us to say (as we do) that we really know something to be true. But nothing I have said need lead us to this conclusion. For that would be to confuse knowledge with absolute certainty.

Obviously, we often do know that something is true, and know it in such a way that doubt becomes an unreal option. When, for example, I return in the evening to my house, enter this room and sit at the desk I am now sitting at, it seems unrealistic to doubt that the desk I am sitting at is my desk. Now it seems that here I am saying that what I see, feel, etc., combined with my ordinary belief in the accuracy of my recollections of my usual surroundings and in the stability of the material world and its contents entails the certain truth of this judgement, if anything is certain. In this case it seems absurd to suggest that the truth is unattainable. But, as Wittgenstein has pointed out, "Here the form 'I thought I knew' is being overlooked".* My feeling of certainty here is due to the fact that I cannot in the normal course of things conceive of any observation I could make, or any modification of the types of justification for my judgements which could come to be accepted, that would lead to my revising my judgement that this is indeed my desk, the very one I was sitting at yesterday. (I am, of course, setting aside the possibility that someone has secretly exchanged my desk for another just like it. Though this is possible, it is not the sort of doubt I am thinking of here.) The fact that something is inconceivable in practice does not make it impossible in principle. However, since there is a corpus of ways of making judgements which is so intimately bound up with

* L. Wittgenstein, 'On Certainty', ss.21

the type of creatures that we are and our entire understanding of the type of world we inhabit that there is no significant possibility of a change in what will count for us as an adequate method of judging, we would expect that the idea of what counts as judging objectively here, and hence of what is true, to remain constant. Thus we can speak with assurance of knowing the truth in this and very many similar cases. This is still the-truth-as-we-believe-it-to-be, though, and not necessary truth. The possibility that we might be mistaken always remains, however far in the background, even though so radical a reappraisal of the grounds we have for almost every judgement we make about the world is not to be expected. Yet if the possibility remains, we must resist the temptation to elevate what is for us inconceivable to the level of a logical impossibility, and treat that as a general principle. If we do, we will have to accept that nothing that counts for us as an objective method of judgement could possibly be revised or superseded. And we would hardly wish to admit that, for the reasons which have been set out in this chapter.

So far I have discussed both the distinction and the interaction between objectivity and truth mainly in terms of empirical judgements, where the need to make the distinction is greatest because the two concepts are at their closest. However, the original reason for making this distinction was to discredit the argument that moral judgements could not be objective because there are no such things as moral truths. That argument has certainly been weakened by what I have said so far, because the non-identity of truth and objectivity makes it impossible to assert that merely because we do not call moral judgements true we can infer that moral judgements are not objective. However, my analysis so far has maintained that there is a relationship between objectivity and truth, even if it is not one of identity or equivalence. Without the involvement of the concept of truth in moral judgements, could we nevertheless talk about there being moral objectivity? I believe that we could do this if we needed to, although it would have the consequence of

making us treat moral thought as something which is in a very early stage of its evolution.

In Chapter One, I argued against the suggestion implied at one point by Westermarck that either the truth or the objectivity of a judgement requires the existence of something real in virtue of which the judgement is true or objective as the case may be. Now the danger in refusing to adopt a Realist position is that the only alternative would seem to be some kind of Conventionalism. This would reduce objectivity to a consensus notion, by suggesting that people are free to adopt whatever method of judgement they like and just call it objective because it is widely held. But this in turn would rob objectivity of its significance, by removing the requirement that anyone has to judge in any particular way because not to do so would be erroneous or irrational. A very similar problem arises in the Philosophy of Mathematics, where the dilemma is between Platonism and Constructivism (which is one form of Conventionalism). The Platonist holds that mathematical truths are truths in virtue of the real relations that exist between mathematical objects; while the Constructivist, unconvinced by references to such real objects or to the real relations that are supposed to hold between them, maintains instead that mathematical truths are no more than truths by convention, and admits the possibility that we might abandon our existing conventions in favour of others — conventions, say, which prove more useful in the light of experience. There is no need to dwell here on the difficulties inherent in both these positions. We need only note that in response to them some philosophers of mathematics have sought a middle way by arguing for a view of mathematical necessity based on meaning. Michael Dummett, for example, has suggested that the meanings which mathematical terms have acquired impose assertability-conditions on us in our calculations, and so make us recognise the necessity of mathematical proofs as they come to be discovered. He writes: "We ought to interpose between the Platonist and the Constructivist picture an intermediate picture, say of objects springing into being in response to our

probing. We do not make the objects but must accept them as we find them (this corresponds to the proof imposing itself on us); but they were not already there for our statements to be true or false of before we carried out the investigations which brought them into being. This of course is intended only as a picture..."*

The great advantage of basing necessity on meaning is that meaning is not determined by individuals deciding what something shall mean, but nor is it determined by the truth of anything independent of human thought. It is thus in a sense conventional, but the conventions cannot be regarded as arbitrary, and are not subject to arbitrary change. As will be apparent from the later development of this thesis, I have considerable sympathy with the general notion of meaning providing a non-Realist but also non-arbitrary basis for judgement, although I suspect that Dummett is wrong if he is supposing that assertability-conditions could entirely replace truth-conditions in mathematics. The value of assertability-conditions is that they would allow mathematics to get started, and so begin to generate conclusions which we can then recognise as mathematical truths. These truths do not make the axioms and principles of inference by which we arrived at them valid, but in the light of these truths we can see that they are. The axioms and principles thus provide justified methods of judgement, and the truths provide a standard against which the results of subsequent acts of judgement by these methods will stand or fall. The objectivity of mathematical proof will then be assured as a route to truth and itself capable of modification and limitation in the light of the truth. (It seems to me that a greater awareness of the interplay between objectivity and truth would make the development of this aspect of the Philosophy of Mathematics rather easier.)

A similar account could be given of the relationship

* 'Wittgenstein's Philosophy of Mathematics' in 'Wittgenstein: the Philosophical Investigations' ed. G. Pitcher, Macmillan 1968, p.447

between truth and objectivity in Ethics, but only provided that we are prepared to regard Ethics as an as yet underdeveloped field of knowledge, and in that respect different from mathematics. Then a truth-free but objective Ethics would be conceivable if we were to regard the objectivity of moral judgement as still firmly entrenched in assertability-conditions derived from the meanings of the terms employed in making such judgements. For morality might not, or might not yet, have generated a body of truths sufficiently clear, or sufficiently well recognised as truths, for them to be available to act as an independent check on the justifications we have for judging objectively, so as to show up that objectivity clearly for what it is, and so as to allow us to modify or limit those objective methods of judgement in the way I have described in connection with empirical judgements. Mathematics, on the other hand, has clearly progressed to this stage in its development. And that might explain why we easily accept the idea that there are mathematical truths, but are far less willing to speak of there being any moral truths. So, if it is felt that talk of objectivity is impossible except in circumstances where talk of truth is not ruled out, we could still talk of moral objectivity as long as we treat morality as a field of knowledge which is at an epistemically primitive stage of its development.

However, though we could continue to link the idea of moral objectivity to the possibility of coming to discover moral truths, eventually, in this way, there is no real need to do that, just because there is no real need to regard objectivity as inevitably linked to truth, providing only that we can link it to some analogue of truth which performs the same function of being that to which objective moral judgements would provide a route. We are likely to go on linking objectivity to truth itself just as long as we treat empirical judgement, not as a clear example of a type of judgement capable of being objective, but as the only type of judgement which could possibly admit of objectivity. There is undoubtedly a connection between truth and the objectivity of empirical

judgements, just because the expression of the truth is the whole point of making empirical judgements, so that, whatever the relationship between truth and objectivity here, truth must be the end that objectivity serves. That was what was meant earlier by describing objectivity as the means to the truth in cases where it is appropriate to speak of there being truth. But if we turn our attention to non-empirical judgements, it becomes obvious that the point of these judgements is not to express how things are at all, but to express how things are properly to be related to one another, or conceived of, or understood, or done. And <u>there</u> the possibility of there being sufficiently justified methods of judgement (and thus objective judgement) will have <u>its</u> point in the idea that there could be an appropriate notion of rightness or correctness which applies to such judgements. In calling a mathematical calculation correct or saying that a decision about how to act has been made aright, we are not of course implying the existence of anything independent of human thought in virtue of which rightness or correctness can be attributed to these judgements. But neither are we implying that to judge aright or to calculate correctly is simply a matter of following an arbitrary rule or adopting a merely conventional method of judgement. Rightness and correctness imply the existence of a standard which has to be met as well as a method which has to be followed. So in morals, for instance, we could imagine objective methods of judgement normally (but perhaps not always) leading to the making of judgements which were actually right. The knowledge so gained might then sometimes lead us to modify or recognise limitations on some of our methods of moral judgement. And here we have an analogue of the interaction between objectivity and truth which we found among empirical judgements.

If this view is substantially correct, then we need not link the possibility of moral objectivity with the need for there to be any moral truths at all; and we can explore the possibility of moral objectivity without being bound to feel that moral knowledge must be something crude and underdeveloped in comparison with other forms of know-

ledge. There are, I remarked earlier, areas of morality where it is not altogether inappropriate to think of there being truth, because there 'truth' is being used rather loosely to indicate the uncompromising presence of what is really truth's analogue in moral matters, right judgement. Thus, for example, I do not think it is wholly misleading to say that 'the relief of human suffering is good' is <u>true</u>; although it <u>is</u> misleading to say this if it suggests that a judgement of this sort is true in the way that empirical judgements are true, so that it could be disproved by showing that it could not be empirically verified. But the fact that a moral judgement cannot be empirically verified does not mean that a man cannot have any justification for so judging, or that the justification might not count as <u>sufficient</u> justification, making the judgement an objective one. It might be felt that to call any moral judgement true would be to employ the concept of truth improperly. But now this is a matter of minor importance. For if we want to say that the only synthetic judgements which can be called true are those where empirical verification is available, we can also say that such empirical truth is only one of the ends that objectivity is a means to. This being so, we are not disbarred from claiming that moral judgements might be objective even where we cannot, or do not, refer to such judgements as true. And this is a reflection of the different roles that the concepts of truth and objectivity play in our thinking.

CHAPTER THREE

JUDGEMENTS ABOUT COLOURS

So far we have looked, in general terms, at some of the confusions which have traditionally surrounded the concept of objectivity, and in the process gained some insight into the conditions a judgement must satisfy if it is to count as objective. In this chapter I shall concentrate on one particular type of judgement which is capable of being made objectively, and discuss in some detail the conditions which do in fact make such judgements objective. I mentioned in Chapter One that the objectivity of judgements is no less in need of explanation when those judgements are empirical ones than when the judgements in question belong to a category where objectivity is philosophically contentious. Partly for that reason, I shall concentrate on judgements about colours. Colour judgements have traditionally been taken to be a prime example of judgements which describe the individual's own experiences. Equally, though, they are obviously capable of referring to objective features of the material world, however much traditional epistemologists may have doubted whether this could really be so. Colours may also be taken as representative of sensible qualities in general, because the problem of the status of colour judgements is the same in essence as that about any other sensible quality which may be predicated of public objects, and either predicated correctly or incorrectly.

In describing how we do in fact make objective judgements about colour, I am not of course arguing that just because this is how we do judge, this is how we are justified in judging. The descriptive process has some value in its own right, since it will lead us to notice distinctions which traditional accounts of the basis of knowledge often overlook, and reminds us of a point made in Chapter Two, that an objective judgement need not refer to the way things actually are, but that appearances can be objective too. However, the purpose of the description runs deeper than that. For in order to begin to explain

why the way that people generally do judge is the way in which it is right to judge, we must begin to loosen the grip of a number of patterns of thought which maintain a seductive hold on the way we approach this question. And one of these is the notion that we could describe our individual experiences in terms that do not extend beyond the experiences themselves. A closer examination of the way in which we actually do make judgements helps to free us from the grip of such an idea. That idea is one aspect of the pervasive belief that the individual's experiences are not only the source or starting-point of our knowledge, but must also be the logical basis of it. This chapter and the next two chapters will expose the real basis for objective judgement by undermining this view altogether.

To talk about the colour of an object is to say something about one aspect of its appearance. Indeed, a sensation of colour has often been regarded as a prime example of a pure phenomenal occurrence, and consequently discussion of colour has figured largely in subjectivist and phenomenological theories of knowledge. However, the concept of colour which we ordinarily employ is not that of sensible qualities which impinge immediately on our consciousness, but of relatively persisting features of public objects available to our sense of vision only under certain conditions, the most obvious of which is, of course, the presence of light. This is not to deny that there are such things as coloured hallucinations, dreams, after-images and the like, to which there need be no corresponding objects. But it is to deny that when we think of colour, we think primarily of <u>those</u> experiences, rather than the experiences associated with seeing coloured external objects. There are good reasons for this being so, connected with what it means to possess colour concepts, and later I shall examine this point at some length, when it becomes central to my argument. For present purposes, however, we need only notice how we do treat colours as properties of public objects, and not personal experiences only, as we do pains. (It is interesting to note that Aldous Huxley became extremely conscious of this fact under the influence of mescalin. The intense

colours he experienced were intensely of objects. The experiences of colour when he shut his eyes were, in contrast, "curiously unrewarding". The idea, then, that interfering with normal perception is likely to detatch properties such as colours from their conventional material associations, was not what he found.*)

Now because colours are ordinarily regarded as objective properties of things in the external world, we find it natural to talk about the colour a particular thing, or part of a thing, is. We recognise that the thing in question — a yellow book, for instance — may present a different appearance with respect to its colour if the conditions of lighting change sufficiently, but we also recognise that its colour has not really changed, and part of the business of assessing what colour a thing is depends on taking into account the lighting conditions obtaining at the time of observation. To simplify matters somewhat, we can say that the colour of an object is available to our eyes when and only when the light is good. As a rule, what counts as good light is strong but diffused daylight. There is, to be sure, a certain amount of latitude in this requirement, simply because many sorts of natural lighting conditions and some artificial lighting conditions do not differ very much from strong diffused daylight so far as the identification of colours goes. At least, these conditions are similar enough not to interfere with the fairly broad range of hues covered by the most commonly-used colour words. Roughly speaking, then, the yellow book that I bought at the open-air market stall still looks yellow when I pick it up to read by artificial light in the evening. It will not, however, look quite the same shade of yellow that it did earlier; and if I look at the book in very intense yellow light (either natural or artificial) it will seem whiter in hue than it really is, whilst in very dim light such colour as the book then shows will be different again. So in order to speak of the colour (the real colour) of an object such as a book, account must be taken of the state of the light. I shall call those conditions of the light in which an

* see 'The Doors of Perception' (Penguin, 1959) pp.30-38

object presents its real colour, standard conditions of observation, or standard conditions for short. I do not mean to suggest that the state of the light is by any means the only impersonal factor that can affect colour judgement. Obviously, the object being viewed must be neither too far from one's eyes for its colour to be distinct, nor too close to them for the light to be able to fall on it adequately. Some conditions for accurate colour observation are quite recondite. For example, it has been pointed out to me that in painting colour effects always depend to some extent on the relation between juxtaposed pigments, so that an area of paint mixed from equal quantities of black and white medium looks distinctly grey when surrounded with a black border, but distinctly blue when surrounded with an orange border painted over the black.* The ordinary person is not aware of this, any more than the person who does not recognise the Müller-Lyer diagram as presenting an illusion is aware that the lines in the diagram are really the same length. So it may be very difficult to pin down the precise real colour of something in many cases. However, this does not prove, as has sometimes been suggested, that there is no such thing as seeing the real colour of something. It only shows that the conditions in which this can be said to be possible can be stated with more and more precision. And this in turn would not be possible unless those conditions constituted part of the rules for correct judgement, which we have seen it is the function of objective methods of judgement to embody. There is no reason why we should not say, within a tolerable degree of accuracy, that those conditions will constitute the conditions under which the colours we see are the real colours of things.

As a rule, then, we would say in ordinary life that if I place an object a suitable distance from my eyes and look at it in ordinarily good light, then I can see what colour it really is. Now this involves another type of assumption on my part, the assumption that my colour

* I am indebted to Fred Cuming R.A. for this example

vision is not defective. One discovers if one's colour vision is defective by directly or indirectly comparing one's own powers of colour discrimination with that of other people. Thus a colour-blind person is one who, basically, cannot make the colour discriminations between red and green that ordinary people can make. Certain drugs and illnesses can also affect colour vision, but here the patient can usually discover the effects for himself by realising that objects which he knows to be certain colours now appear different in standard conditions. Most people are said to have minor peculiarities of colour discrimination at one point or another in the visible spectrum, but these differences usually cause no problems, because the commoner colour words all cover a broad enough section of the spectrum to mask them. So, in line with the definition of a standard judge offered in Chapter Two, we might say that a standard colour-judge will be one who makes the same colour discriminations that anyone else could make. Probably few of us are actually standard colour-judges over the whole range of the spectrum, but probably most of us are over most of it; and the notion of a standard colour-judge, though imprecise, seems straightfoward enough.

My account, therefore, of the way in which we ordinarily talk about the real colours of things amounts to this: we have a notion of objects <u>being</u> certain colours as well as <u>looking</u> certain colours; and the colours that objects <u>are</u> are the colours they look when seen by standard judges in standard conditions of observation. So when I say that the daffodils outside the window <u>are</u> yellow, what I am saying amounts to the claim that the daffodils look yellow to me, and that because I am a standard judge and these are standard conditions of observation, I am entitled to assert that yellow is the real colour of the daffodils. The connection here with the objectivity of colour judgements is an obvious one. For here I am pointing to features about my judgement that justify me in saying what I do about the daffodils themselves, as opposed to my or anyone else's impression of them. Anyone who judges differently from me here is in error unless he can

show that my beliefs about the current conditions of observation and/or my status as a standard judge are mistaken.*

Of course, my judgement about the colour of the daffodils is, if I am correct about myself and the conditions, not only objective but also true. This should hardly surprise us; for the judgement is an empirical one, and as we saw in Chapter Two, objectivity is the route to the truth in such cases.

(All this might well seem philosophically uncontentious, even banal; but it can lead to the espousal of a particular theory about the meaning of statements about colour which we must be on our guard against. I am <u>not</u> saying here that the meaning of 'X is red' is 'X looks red to standard observers in standard conditions of observation'. If '...is red' were to be defined by reference to '...looks red to...', I should be guilty of two related errors. First, I should be covertly embracing a consensus view of correct judgement; second, I should have to explain how looking red could be understood without reference to being red, although knowing what it is for something to look red presupposes knowing what it is for something to be red. I am not, then, defining '...is red' in terms of how things happen to look to certain people in certain conditions. But it is obvious that a perceptual concept like redness cannot be understood at all without reference to <u>some</u> appearances of things as red; and as it happens the distinction between being and seeming <u>is</u> drawn, in this and many other cases, in line with what standard judges do say in standard conditions of observation. My pointing out that this is so is not to be mistaken for an argument to that effect. The question of why standard judges are justified in judging as they do will be dealt with later on.)

In addition to talking about the colours things are, we may of course talk about the colours things look. How things look depends not only on the colours the things are

* unless the case is a borderline one, when the disagreement may be apparent rather than real. See my 'Vagueness and Colour Predicates' (Mind 1972) pp. 576-7

but also on the conditions under which we come to observe them. Now if the justification for making an objective judgement about the colour something is depends on the claim that one is a standard judge judging in standard conditions, it is equally the case that the judgement that something looks a certain colour can be an objective judgement as long as it also makes reference to standards of this sort. For my claim about the colour something looks simpliciter is not only a claim about the colour it looks to me. It is the claim that any person who qualifies as a standard observer and who is in the same, or roughly the same, situation as myself in relation to the object in question will see the object as looking the same colour as it looks to me. It is often supposed that there is something necessarily subjective about any appearance-statement, but this is not the case. The apparent colour of things changes in a consistent way according to changes in the strength and quality of the light, the relative distance of the object from the observer, and the medium through which viewing takes place. Thus reds change to look like browns as the light becomes weaker, mountains appear bluer in the distance than they do close to, and smoke, mist, glass and the philosopher's coloured spectacles all produce variations in perceived colour which are consistent and predictable. There are very many such deviations from standard conditions of observation which we are all accustomed to, and which lead us to talk of the apparent colours of things in those conditions rather than about the colours things actually are. So it is still perfectly possible to speak of the proper colour that a thing of a particular sort should look in a given set of non-standard conditions when seen by a standard judge. And to make a judgement about how something can properly be said to look in certain conditions is clearly not to make a subjective judgement about the impression one has as an individual, but to make a judgement about an aspect of things which is open to public agreement or disagreement. In short, provided it is made with sufficient justification, such a judgement will count as an objective one. And since I

am suggesting that what standard judges actually do say is indicative of there being, as a rule, sufficient justification for the judgements they make, it follows that if standard judges do say such things as 'yellow objects look white when seen in sodium vapour light', such judgements will in fact be objective. Being empirical judgements, the fact that they are objective will normally mean that they are true, since here objectivity is a means to the truth. The fact that such statements as the above are true is most easily explained by recognising that they are objective, even though they only relate to appearances, and not to how things really are.

But although judgements about the proper look of things are objective when made by standard observers, it does not follow that they are true when considered as judgements about how things really are. If one says that an object which appears white in non-standard lighting conditions is white, what one says may be false, for yellow objects can also look white in some non-standard lighting conditions. And if one says of an object which looks white under sodium vapour light, but which looks yellow in daylight, that it is really white, one's statement is simply false. It is from cases of this sort that we have already been able to see how large numbers of judgements can be objective without being true of their object as it really is. For judgements about colours, it is a consequence of the view that colours are properties of objects, but properties that are available to us only in certain conditions, including conditions about our own competence to judge. Although this view of colours is contrary to that which has most commonly been held by philosophers in the past, it is nevertheless the view of colours which we actually have. I shall argue in due course that a subjectivist view of colour cannot match up to the things that we ordinarily say about colours, because it cannot really amount to a theory about colour at all.

I have now suggested (though not yet attempted to prove) that when a colour judgement is made by a standard judge, the fact that he is a standard judge is a reason

for saying that as a rule his judgement will be a sufficiently justified one, and thus objective. Judgements made by standard judges, both in standard and non-standard conditions of observation, then, are in their different ways objective. Now clearly reference to standard judges is reference to anyone who happens to qualify as one, so the judgements I have been discussing must be such that anyone could make them. Further, we expect standard judges to agree in their judgements. The judgements we have been considering are obviously public ones. Both forms of objective judgement are to be contrasted with the individual's immediate awareness of colour per se; the perception of a sense-datum of (a) colour, as it is sometimes put. It is no doubt odd or artificial to try to describe an experience of colour without any reference to coloured objects in the public world, even when the colour is not attributed to anything in that world. If, for instance, I say 'there seems to be something red before my eyes', I am expressing, in a hesitant way, an experience that I do not quite know how to characterise but which I believe to be in some sense comparable to there really being something red before me. If I say 'when I shut my eyes, I am conscious of a sort of red haze', then again the experience is being expressed in terms of its being like a real red haze, one that standard judges would agree either was red or really looked red.

But for all that, the epistemologist is quite entitled to introduce descriptions of events which we would not normally describe in that way, in order to produce his rational reconstruction of the structure of knowledge. He is entitled to invent a neutral description for the having of a sensation which involves no claim whatever about the external world or its appearances, providing, of course, the result of his doing so is a theory which explains, or helps to explain, our ability to make the judgements we can in fact make. So we may explore the possibilities opened up by the philosopher's locutions 'I have a red sense-datum' or 'I sense a redness', as long as these expressions help to clarify the sense of what we really do say about

colours by giving coherent expression to the subjective aspect of our experience.

Unfortunately, this is just what such neutral expressions about colour phenomena fail to do. They do not do it because they cannot do it. For they contain colour words such as 'red' which are part of the ordinary language we use to characterise the colours which public objects are seen to have by us all. And if words like 'red' are to be used in some way other than this, their connection with the normal usage of colour words must be made clear. Now there only two possibilities. Either colour words are used in reports of sense-data in just the way they are ordinarily used, or they are not. (If they are not, their obvious point of application is the particular phenomenal quality of the sense-datum which the individual has.) Let us consider the first possibility. Here the words used by the individual to describe his immediate awareness of colour are being used in just the same way that standard judges use them to describe public objects. This means that whenever the individual is confronted by something which standard judges would describe as being red, his report of the sense-datum he has will also contain the word 'red', unless circumstances which standard judges would recognise as having a distorting effect on colour judgement are affecting the individual in question, in which case he will say something different. In other words, the individual uses the word 'red' exactly as we all do, except that in his case the word is supposed to refer to the quality of an experience itself and not to the real or apparent property of an object which is the subject of the experience. In one way, then, all this talk of sense-data is a harmless irrelevance. But suppose the sense-datum theorist says that it is not, and that reference to sense-data marks an important difference. In that case, we are entitled to ask what difference it makes, and in particular why it should be that the sense-datum expression and the ordinary way of talking are the same in every case. Coincidence being too far-fetched, there must be a sense in which the individual's experience is comparable to that of standard

judges. Now one person cannot have another person's phenomenal experiences, so neither he nor we can compare our individual's phenomenal experiences directly with those of standard judges. All comparison must take place by reference to what is publicly available. So reference to objects is at some point unavoidable. The individual who calls his colour experience an experience of phenomenal redness must, therefore, at some point be able to refer to what others see as red, and in virtue of the fact that they also see it as red. He must recognise that his characterisation of a feature of his experience as being an experience of seeing <u>red</u> involves what is qualitatively like what he sees when he sees things that really are red — that is, things that look red to standard judges in standard conditions.

But in that case, the claim that locutions such as 'I have a red sense-datum' can be used in such a way as to avoid reference to any publicly available object, is true only in a very restricted sense. In a particular case, the description of a hallucination, for instance, the expression might be used to describe a phenomenon where there was no correspondingly propertied object available to sense there and then in the public world, or where the question of whether there was such an object remained open for some reason. However, a language consisting of such expressions could not be neutral with respect to the existence of <u>all</u> publicly available objects and their properties; for the justification for the use of the words would depend on there being <u>some</u> points of comparison between how the individual describes his experiences (even to himself) and what standard judges would say about public objects. So an account of phenomenal experience which is presented in this way presupposes the ability to refer to at least some public objects and their publicly observable properties of the requisite kind. Words such as 'red' could not be used in just the same way in a 'phenomena-language' and our ordinary way of speaking of things without the phenomena-language presupposing our ordinary way of speaking about things. In that case, words like 'red' could not be used wholly without reference to public objects, as the suggestion

presently under consideration claims they could be.

It might perhaps be thought that the dependence could be the other way round; that is, that words like 'red' could not be used in the ordinary way about public objects and their properties without that presupposing a phenomena-language on which ordinary language is based. This was the assumption underlying classical Empiricism. The wrongness of that view will be extensively treated of in the next two chapters. In essence its mistake lies in asuuming that it is possible to talk in the other way I mentioned, by saying that there could be expressions which refer only to the individual's own experiences and not to anything else, so that in the first instance they would have no connection with any public usage, because public usage was to be constructed out of them. But assuming, for the moment, that expressions for such logically private experiences could occur and be meaningfully used by the individual, there would still be the problem of how such words could ever form the basis of a public language for objects and their publicly-observable properties, or else how the individual's language could ever connect with any such public language. In either case, it would be something more than a happy coincidence if two individuals used the same term at the same time and meant the same by it. This is because two uses of a word whose function is referential cannot constitute two uses of the same word (i.e. not just two uses of the same symbol, but two uses of the same symbol with the same meaning) when there exists no possibility of their being used with the same reference. And this is excluded ex hypothesi.

There is a powerful argument, put foward by Wittgenstein, to the effect that a word used to describe a logically private experience, as is the case with 'red' in the theory being considered here, would not be a word at all. That is to say, it would not have any meaning, even for the individual who used it solely to characterise his own experiences to himself.* I shall refer to this argument in detail in Chapter Five, where it becomes

* 'Philosophical Investigations' I, esp. ss. 258-279

central to my argument. I need not rely upon it now because more is proved there than is needed at present. Here it is sufficient to point out that even if the word 'red' could be said to have a meaning if that meaning were restricted to describing the private experience of an individual, it would not be a word capable of being applied to public objects because it would not be a word about the application of which it would be possible for standard judges to agree. This is because, ex hypothesi, the use of the word is restricted for any user to his own experience alone. With this restriction in force, it would be impossible to speak, as we do, either of the colours things are, or of the colours things really appear to be. The price of accepting the meaningfulness of private colour-words would be the abandonment as meaningless of the descriptions we give of one important feature of the world of public things. Of course, this argument applies to much more than colour-words alone. All sensible qualities which were defined in the same way would be rendered unavailable for use in describing public objects. Needless to say, this is too high a price for us to pay for the freedom to describe sense-data without reference to a public world.

It has been suggested that there is an intermediate position between the two I have just been discussing, viz. the use of colour-words to refer to private phenomena but in a way consistent with what standard judges say, and on the other hand a use of colour-words to refer to private phenomena in a way that takes no account of what others may say. A.J.Ayer has discussed such a position.* He imagines a Robinson Crusoe figure, alone from birth, who has come to invent a language for himself. The language includes words for his sensations, including words for his colour sensations no doubt, although Ayer concentrates upon words for inner feelings. These feelings are not, of course, publicly accessible, but Ayer accepts that they might be revealed by being associated with "natural expressions" through some form of characteristic behaviour.

* 'Can there be a Private Language?' (PAS 1954) repr. in 'Wittgenstein' ed. G.Pitcher

However, what Ayer says ought to apply equally to words for colours, except that they, presumably, will not be associated with any obvious natural expressions. Ayer's castaway is now joined by a Man Friday, and according to Ayer, Crusoe will be able to teach Friday the names of objects, although he cannot teach him the names of sensations which are "entirely private", i.e. which "have no 'natural expressions'" associated with them. This would presumably include colours. Yet if the resulting language is to be at all like our language, there must be not only names for objects but descriptions of objects as well. For example, Friday must be able to discriminate between two otherwise identical goats if Crusoe says 'bring the white goat', so colour descriptions ought to play a part in this language. Now we can imagine Crusoe teaching Friday the meaning of 'white goat' by ostension perhaps, though this is not the same thing as teaching Friday what 'white' means. 'White goat' would merely identify the animal, not describe it, and we can imagine Friday using the word 'white' in a number of cases without having any inkling that white is a colour, unless of course Friday already has a language of his own which already allows him to make sense of universals, so that his problem is one of translation. (But translation from what? It is a real, though separate, problem for Ayer to explain how Crusoe comes by a satisfactory understanding of universals. For colours are at the moment names for occurrences.)

At any rate, according to Ayer, a word might serve initially to refer to a private phenomenon of the sort described earlier, and then have its meaning extended, as it were, from the private to the public realm, by becoming applicable to objects rather than sensations. We are to imagine this happening without the words Crusoe uses changing their meaning for him. But this is impossible. For if the colour words did originally designate wholly phenomenal features, they could not be extended in this way, as I have argued in the previous few pages. On the other hand, if by some mysterious means (for ostension cannot be the whole story) Crusoe can teach Friday the names of colours as properties of objects, then the names

of colours must have been applicable to objects from the first. They have not acquired an extended meaning with the arrival of Friday. They must always have been applicable to objects which could have been described by others, had there been others, using the same descriptive terms, so that is what they must have meant to Crusoe when he invented them. In this case the claim that colour-words are part of a language which describes only the phenomena of Crusoe's experience is false, as my first argument showed.

The heart of our difficulty with Ayer's account lies in understanding how, from the beginning, Crusoe could be said to be attaching words to sensations rather than to what he saw. In Crusoe's world it is not at all clear whether this distinction can even be made. Ayer himself expresses some doubt on this point, saying that Crusoe is not bound to make this distinction, and that he is not bound to accept the same criteria of identity for an external object that we do. He says only that "it is reasonable to suppose that they will be the same" without saying why this should be reasonable. I assume that Ayer just supposes that Crusoe's experience will seem to him like ours does to us. But if my argument is correct so far, he can have no way of finding that out (if indeed it means anything to talk about one man's private experiences being the same as anyone else's without some reference to a shared system of concepts and judgements). Ayer's supposal here seems charitable rather than reasonable. At any rate, the basis for drawing the distinction which Ayer offers is the variability between different sensations, some of which will, it is implied, suggest the existence of an external world to Crusoe. Ayer calls this "the principal distinction which he is likely to draw between 'external' objects and his 'inner' experiences". However, apart from the possibility of Crusoe's constructing a purely theoretical (and wholly untestable) causal theory of sensation which might involve the postulation of something called 'an external world' as a causal basis for sensation, there seems no real likelihood that it would ever occur to Crusoe to draw such a distinction. For all he has to suggest such

a distinction to him is the qualitative differences within the range of his phenomenal experiences. The very best Crusoe might be able to do would be to label his most regular experiences 'external experiences' and postulate 'objects' for them to be experiences of. According to Ayer, Crusoe will recognise that "his experiences are transient in a way that external objects are not". This will be true by definition for Crusoe, though we will have no idea what he means by 'object' and 'experience'. But in any case, this will not give us the distinction we have. For the distinction between the inner world and the external world cannot be based on the fact that some experiences last longer than others, or that they are more or less frequent than others, or more regular, or more intense. The distinction does not depend on qualitative differences between sensations at all, but on a difference between the circumstances which surround the characterisation of experiences as being of one sort or another. In fact, it depends on whether the judgement of what the experience is fulfils the conditions for being an objective one. A dream or a hallucination or a memory of a rainbow need not be different in quality from what we call seeing a real rainbow. What separates the two kinds of experiences is that others can see the real rainbow as well as myself. The main thrust of Ayer's whole argument against Wittgenstein is that the understanding of all experience, whether public or private, must ultimately depend on the individual, for it is the individual who must recognise that experience for what it is. But that is beside the point so far as we are concerned. The very possibility of a distinction between inner and outer experience rests upon others being able to share my experiences, irrespective of what it means to be able to recognise an experience. Crusoe cannot make this distinction without the possibility of there being a contrast between the experiences he shares with others and those he does not. This possibility exists whether or not Friday has appeared on the scene, though we could understand Crusoe not realising this until Friday was there and shared Crusoe's language. But if the possibility is ever to be there, it must always have been. (The question of whether

Crusoe could ever have invented a language for himself, or whether what he invented would really be a language, is a separate issue, which I shall discuss elsewhere; although the belief that an individual could invent a private language to characterise his phenomenal experiences is an integral part of the confusion involved in the belief that language might refer first to phenomena, and later become extended to take in references to objects with independent existences.)

All this must lead us to reject the usefulness of trying to give an existentially neutral account of colour phenomena in an attempt to understand the logic of our discourse concerning colours. A report of seeing a colour must, as I have argued, be couched in language which allows for the possibility of others having that experience too. It does not follow that such a report must be objective, though. The person having the experience may realise, or suspect, that he is having a different experience from that which a standard judge would be likely to have in his place. This leads us to see the point of 'looks-to-me' statements. Such statements have two related functions. First, they state the nature of the phenomenal experience in terms of its similarity to the observations of standard judges (which, of course, implies that the person having the experience is in some cases at least himself a standard judge). An expression such as 'this book looks yellow to me' therefore reports an experience comparable to the experiences of standard judges in standard conditions when they look at something which is yellow. Secondly, the expression indicates caution about what standard judges would say in this case, and so invites confirmation or denial by other observers about the colour the book really is, or about its real appearance. That is the qualifying function of the words 'to me'. It therefore asserts much less than the standard judge's expressions 'this book is yellow' and 'this book looks yellow'; but it does not report a different kind of event. It reports an experience (but not an inner experience) and attempts to relate it to how things are or appear in the world, and may indeed succeed in doing so. For if the judgement is in fact a justified one (as

would be indicated if it turned out that standard judges agreed with it, so that I could on this occasion count myself as a standard judge too, in spite of whatever consideration made me hesitate to believe that I was) then the book which looks yellow to me really does look yellow, and my judgement is an objective one. If, on the other hand, standard judges contradict me, for the good reasons that make them count as standard judges, then my experience is a subjective impression rather than anything which approaches the truth of the matter. I now realise (providing I accept their judgements as standard ones) that the book which looks yellow to me only seems to be yellow, but neither really is it or really looks it. Neither yellowness nor 'yellowlookingness' is therefore one of the book's objective characteristics — that is, it is not one of the book's characteristics at all. It is instead a feature of my own experience in seeing the book and nothing more than that, and it would be a mistake to predicate anything of the book on that account. (This is how, I suspect, we come to understand the notion of phenomenal experience, as a contrast developed out of our notion of experience simpliciter, which is primarily of public objects. And this is why it is the subjective and not the objective statement which needs additional words to qualify it. I shall refer to this point again in Chapter Five.)

I have now discussed three types of judgement which may be made about colours — the colours things are, the colours things may appear to be, and my individual impressions of colour — and tried to show how all three types of judgement embody reference to the public properties of objects in the external world for anyone who has the concept of colour that we have, as is revealed by the things we actually say about colours. By all being related to the public properties of objects, these three types of judgement are all related to one another. This relationship shows up most clearly in terms of the relative corrigibility of the different types of statement. I have already said something about this earlier in this chapter when discussing the relationship between judgements about the colours things are and the colours things

look, and also in Chapter Two when discussing the objectivity of the different kinds of judgement we might make about the Müller-Lyer Diagram. Now, however, the picture may be surveyed rather more comprehensively. First, it might well seem that one's subjective statements about how things look to one will be automatically correct, so far as they go. (And even this is not quite true. One can sometimes modify one's own subjective impressions, by concentrating harder, for instance. Thus a vague impression of yellowness may give way to a more precise impression of yellowy-greenness. Are there two impressions here, or one impression with two descriptions of it? By insisting on the incorrigibility of sense-impressions, the sense-datum theorist creates a problem for himself here, which disappears in my account because concentrating on the impression one has can be taken as making a move towards a more objective mode of judgement.) This difficulty apart, the virtual impossibility of being wrong about how things look to one, which for various reasons has so impressed many philosophers since Descartes' time, amounts to no more than the truism that one has what one has. Although the experiences expressed by statements about how things look to one may be thought of as being, in Kant's phrase, the beginnings of all our knowledge, we should not confuse the most limited and open form a knowledge-claim can have with what is supposed to be the epistemic basis for any claim to knowledge derived through the senses. A statement about how things look to one is the beginnings of knowledge in the former, but not the latter sense. Judgements about these experiences may not even be about the world at all, except in the attenuated sense in which the form of expression that the judgement takes must admit of the possibility of reference to objects in the world and their properties, in that the terms used must be applicable to things in the world, as already explained. Now these looks-to-me judgements are corrigible by reference to objective judgements about the way things actually appear in the world, but not vice versa. Thus, while it may be the case that when I have a subjective experience of colour,

the experience I have is the experience I have, the question of whether what I see looks, say, yellow is settled not by me but by what standard judges say — except that insofar as I myself am a standard judge, I may be able to decide the question for myself. However, my ability to do this depends on my <u>being</u> a standard judge, and this is only ascertainable by comparing my colour judgements with those of others, through the medium of our common language. If, as I have argued, the way we do in fact use words for colours must involve reference at some point to how things in the world really are, or really appear, then 'looks-to-me' statements, in their primary application, are always capable of being overridden by statements about how things really look. The objective statement always corrects the subjective one.

A similar relationship exists between objective judgements about appearances and objective judgements about appearances in standard conditions — judgements, that is, about how things really are. For an objective judgement made by standard judges in non-standard conditions is corrigible by virtue of the judgements made by standard judges in standard conditions, plus the recognition that in the former case conditions <u>were</u> non-standard whereas in the latter case they are not. Again, the converse of this does not hold; and again, this is in recognition of the fact that certain conditions for judging count as those where judgements are actually capable of being correct. (In the case of judgements about sensible qualities, these conditions would in practice be those in which the largest number of discriminations could consistently be made by the largest number of standard judges.) Thus we can see that the most correct answer to a question like 'what colour are those geraniums?' would <u>not</u> be an answer in terms of how they look to me, now, but an answer in terms of how they would look to standard judges in standard conditions. Similarly, the most correct answer to the question 'what colour is before you now?' would be the one given by a standard judge in standard conditions, because 'colour' is a term for an objective property of things, and the standard judge in standard conditions is the person best placed to give a correct answer.

But it will be objected by the sense-datum theorist that one's own impression, however much it differs from what others might say, is nonetheless in itself not corrigible. By saying that, he shows he has missed the point of what I have been saying. He is confusing his actual impression with his description of it in words. However much it is undeniable that he has the experience he has, his characterisation of it relies on words which we can also use, and his use of those words lays his judgement open to the processes of corrigibility which I have described, and which do in fact operate on those words as we use them. The expression 'I have a red sense-datum' is an attempt to detach the word 'red' from its real setting in our attempts to describe how things really are in a world of public objects and their properties. The notion of a sense-datum as something distinct in itself and something capable of being accurately described gives a spurious credibility to the hesitant, weak, aberrant or uncheckable impressions to which we are all prone, and which we express, when we are being sufficiently guarded, in judgements about how things look to one. And these expressions, as I have said, are corrigible by other expressions which take more account of the requirements a judgement must satisfy if it is to be objective, and so express how things are with a proper degree of correctness.

In concluding this account of how we do in fact make judgements about colours, there is one other essential feature of our capacity to judge which deserves to be mentioned, partly for the sake of completeness, and partly because some reference was made to it in the previous chapter. This is the fact that objective judgements demand the employment of a relevant sense. I have mentioned the fact that light is necessary for colour judgement to occur, but not the obvious fact that eyes are also necessary. No doubt this is a harmless omission; but if it were recognised that there were other ways of making colour discriminations as well as by sight in the presence of light (e.g. by touch), the question of whether the eyes or some other organ of sense were involved in the judgement would become an important one, and highly relevant

to the question whether the judgement was to count as an objective one or not. For example, shapes are generally discernible by sight and by touch, whereas weights cannot be directly detected by eye, although with practice the appearance of an object can often be a reliable basis for estimating its weight. This is clearly an inferential method of judgement; but it is less clear whether the look of a surface is a direct or inferential method of judging its texture. In cases where something looks oval but feels round, or looks heavy but feels light, or looks rough but feels smooth, it is important to know which, if any, of these judgements employ the appropriate methods for objectivity.

But let us imagine a case, like that referred to in Chapter Two, where a person seems to be able to make colour discriminations by touch alone. Would we want to say that what this person was judging with his fingers was colour? To some extent, the answer we give depends on whether we are prepared to adopt a causal theory of colour which does not restrict the word 'colour' to what is detected by any one organ, but is prepared to extend its usage to any method of detection of what is, causally speaking, the same property of objects. But whether we want to call such an ability to discriminate a direct or inferential method of judgement depends on the concept of colour that we have. Now we certainly cannot deny that an ability to detect what we would call colour-differences by touch alone constitutes an inferential visual sense of some sort, as long as the discriminations available to touch are exactly, or almost exactly, the same as those available to standard judges of colour judging by eye alone. This detection of colours by touch is clearly parasitic on detection of colours by sight, just so long as it is an oddity. But if it were to become commonplace, so that the 'feel' of colours became something about which people's tactile perceptions might sometimes differ, but normally would not, then standards of objective judgement would become applicable to 'colour-feels' directly. Assuming that people could go on making the same discriminations by eye alone or by touch alone if they chose to, there seems no reason why we should not call both activities colour per-

ception. If we did that, the concept of colour we should then have would be different from the one we have at present, though the new concept would include everything our present concept contains except the restriction of what can properly be called colour to what can be seen. The point is that if awareness of colours via more than one sense were to become normal, then felt colour and seen colours could equally be the subject of standard judgements about colour, and felt colour would then be as much an objective property of objects as seen colour is.

I cannot imagine what it would be like to feel the colours of things, any more than a congenitally blind person can imagine what it is like to see the colours of things. Yet I do not need to be able to imagine what it is like in order to appreciate what it would mean to say that such judgements were objective. For I know in general terms what it is for a judgement to be objective, since I know what it is for there to be judgements of mine which are made in an appropriate way and which meet the standards of objective judgement. They are the interpersonal standards of proper judgement which people normally agree in, and which lead us to say that, at least in the clearest cases, we are aware of things as they really are.

Now one of the advantages of the account of the different types of judgement I have given so far is that it enables us to see through that venerable fog, the Argument from Illusion, in at least some of its forms. This argument, one of the main props of subjectivist scepticism, is supposed to show that we cannot trust our senses to tell us about the external world (which, of course, we cannot find out about in any other way either). The core of the argument is simply this: we know our senses sometimes deceive us; therefore the senses are inherently untrustworthy. (Usually, however, the formulation of this argument is deceptively complex.) In its simplest form, though, the fallacy is transparent. For how do we know that our senses sometimes deceive us? We can only know this if we know the difference between what counts

as getting some judgement right and getting it wrong. But this is simply the ability to recognise the difference between making a judgement with sufficient justification and making it without such justification — the difference between objective and subjective judgement, that is. And that is something we do know about. So it does not follow that our senses are untrustworthy. For the fact that we can distinguish right cases from wrong cases shows that we know <u>when</u> we can trust our senses and when we cannot — or if and when we do not know this, we know that others will point it out to us, and we will come to recognise our mistakes.

It will be found that any formulation of the Argument from Illusion, no matter how complex, rests on a confusion between the three different sorts of sensible judgement which I have been at pains to differentiate from one another in this chapter. Ayer, for example, presents one form of it, in part, as follows:

"Now considering first the fact that appearances vary, we may argue that this proves at least that people sometimes do not perceive things as they really are. If, to take a familiar example, a coin looks at the same time round to one person and, from a different angle, elliptical to another, it follows that it is to one of them at least presenting a deceptive appearance. The coin may in fact be neither round nor elliptical; it cannot in any case be both. So that if each of these persons judges that he is perceiving the coin as it really is, at least one of them will be undergoing an illusion."*

We can see at first glance that here the argument trades on the possibility of confusion between seeing the coin as it really is (e.g. as round) and seeing the coin as it really should appear from a certain angle (e.g. as looking elliptical). And certainly Ayer is not treating appearance-statements as carefully as he ought to do. It would, for instance, help if it were pointed out that seeing a three-dimensional object other than a sphere as

* A.J.Ayer 'The Problem of Knowledge' (Penguin 1956) p.87

it really is involves taking into account the angle from which it is viewed. Only seeing from certain angles will amount to standard conditions of observation for judgements involving shapes. So that if we do not take factors like this into account, we cannot even begin to talk about seeing things as they really are. If we cannot do this, there is no point in trying to talk about failing to do so sometimes, and the Argument from Illusion cannot even get started in this form.

However, the really interesting point of confusion in the argument we are considering lies in the fact that Ayer regards the alternative to seeing something as it really is (such as seeing a round coin as elliptical, and supposing that it really is elliptical), not as a case of mistaken judgement, but as a case of illusion or deception. J.L.Austin has taken Ayer to task for a similar misuse of words in another work.* Austin reminds us that "it is important to remember that talk of deception only makes sense against a background of general non-deception...It must be possible to recognise a case of deception by checking the odd case against more normal ones."** Austin's point is a sound one, but we can go a little further. Ayer's choice of words is more than just a confusion or a suggestio falsi; we can see that he does not recognise the significance of objective apparent judgements, and really believes that the only alternative to seeing something "as it really is" is to 'see' something which has no basis or counterpart in objective experience at all. This is the real danger, not only of neglecting the possibility of there being objective judgements about appearances, but, more important still, of failing to appreciate that even the most subjective description of a phenomenal experience belongs to a form of discourse which is founded on our capacity to make objective judgements.

To be fair to Ayer, he does not accept the validity

* 'Sense and Sensibilia' ed. G.J.Warnock (OUP 1962). Austin is referring to Ayer's 'The Foundations of Empirical Knowledge'.

**op.cit. p.11

of the Argument from Illusion in the form I have quoted from his writings, although he feels it to have "much persuasive force". But though he has doubts about the Argument, the way in which he presents them subsequently, in the work I have quoted from, shows that he is still in the grip of a sophisticated form of sense-datum theory, which makes immediate appearances to the individual the only possible trustworthy source of knowledge about the external world. Since Ayer knows, and we know, that an appeal to the individual's subjective judgement is inherently untrustworthy, the Argument from Illusion will continue to exercise a seductive sceptical appeal unless the belief that the individual's perceptions by themselves are the rational basis for knowledge and objectivity is thoroughly abandoned.

There are, of course, many other examples of this type of sceptical argument in philosophy textbooks. The reference to one of them in the last few pages might have seemed to be incidental to my main argument, but this is not really so. For apart from emphasising the difficulty of taking a subjectivist approach to the foundations of knowledge, it serves to introduce a rather similar argument about the variability of our experiences, which presents a problem which my argument must eventually overcome. It will be remembered that in Chapter Two I put forward an account of objectivity which rested on the idea that an objective judgement is one which we are sufficiently justified in making. In this chapter, on the other hand, I have tried to describe how we actually do make various types of judgement in practice, and characterised their objectivity by reference to the standards that we actually do employ. It was convenient to do this because it led easily into my criticisms of the notion of neutral sense-data, and so to the idea that references to colours must be primarily to the qualities displayed by publicly-accessible objects, and not to the subjective experiences of individual percipients. But, as I warned earlier in this chapter, it would be wrong to imply that the standards we all happen to apply to judgements, when we judge in what we hope is an objective way, automatically qualify as sufficiently justified forms of judgement just because

we happen to use them. The fact that we all as a rule judge in the same way does not in itself prove that our judgements are right. This is something that Bertrand Russell has drawn attention to in a passage in 'The Problems of Philosophy', where he is employing a variant on the Argument from Illusion to suggest that no experience can be thought of as more real than any other.* His argument runs as follows:

"Although I believe that this table is 'really' of the same colour all over, the parts that reflect the light look much brighter than the other parts, and some parts look white because of reflected light. I know that, if I move, the parts that reflect the light will be different, so the apparent distribution of the colours on the table will change. It follows that if several people are looking at the table at the same moment, no two of them will see exactly the same distribution of colours, because no two can see it from exactly the same point of view, and any change in the point of view makes some change in the way the light is reflected... It is evident...that there is no colour which pre-eminently appears to be <u>the</u> colour of the table — or even of any one particular part of the table — it appears to be of different colours from different points of view, and there is no reason for regarding some of these as more really its colour than others. And we know that even from a given point of view the colour will seem different by artificial light...or to a man wearing blue spectacles... This colour is not something which is inherent in the table, but something dependent on the table and the spectator and the way the light falls on the table. When, in ordinary life, we speak of <u>the</u> colour of the table, we only mean the sort of colour it will seem to have to a normal spectator from an ordinary point of view under usual conditions of light. But the other colours which appear under other conditions have just as good a right to be considered real; and therefore to avoid favoritism, we are compelled to deny that in itself, the table has any

* OUP 1912, pp.8-10

one particular colour."

This way of speaking of colours is open to some of the objections I have already made at some length against treating colours as simultaneously pure phenomena and properties of public objects. Furthermore, Russell applies the word 'real' equally to all manifestations of colour, which deprives the word of its proper role in marking a contrast between some appearances of an object — the standard ones — and others. However, setting aside these difficulties, Russell's formulation of the problem does have the merit of casting suspicion on what he suggests is only the normal way of seeing tables and the like; for by implication he is questioning whether the standards we do in fact adopt as justified in our talk about the real colours of things in the external world are to be preferred to any other standards which might be chosen, perhaps arbitrarily, to fulfil this role.

Now it was explicit in Chapter Two, and has been implicit in this chapter, that our capacity to talk of how things really are, or really look, necessitates reference to some public standard of application, which will sometimes differ from one's own individual judgement, and against which one's own judgement stands liable to correction. If I am right in this, then nothing Russell says about the equal reality of all colour perceptions can stand scrutiny. Yet he does raise the difficult question of why what he calls "only...the sort of colour it will seem to have to a normal spectator from an ordinary point of view under usual conditions of light" should be counted as the standard to adopt. In one way, there is an obvious answer to this by pointing to the most frequent way of seeing things as being, so to speak, the democratic choice for a standard. However, this would be a quite inadequate justification. We have already seen, in Chapter Two, that the standard, whatever it is, must be such that it can on occasion overrule a majority verdict, and show us that the majority can fail to be objective in their judgements. So just saying that the way most people do judge constitutes the standard of justified judgement will not do, no matter how precise we are about how people actually

judge. Even if we say that most people will judge in one particular way for one particular reason, we still cannot regard such a judgement as ipso facto a justified one. The descriptive account of objective colour-judgements given in this chapter can only be proved to be a correct account of objective judgements if it can be shown that there is a correspondence between the way people generally do judge and the way in which it is correct to judge.

Fortunately, the way in which Russell presents his argument gives us a tool with which to work on this question. For in the sentence which I quoted on the previous page Russell vigorously conflates two notions which, if separated, will supply us with the beginnings of an answer. In that sentence, Russell refers to standard judgements indifferently as involving what is ordinary, usual and normal. This is most unwise; for the fact that he treats these terms as equivalent to one another shows that he is unaware of the real issue here. He uses all three terms to refer to features of the most common way of seeing things, and rightly regards that as an inadequate basis in itself for claiming to have judged correctly. However, the normal way of seeing things is not, as Russell thinks, necessarily the most common way of seeing things. The normal may on occasion stand in contrast to the usual, though it could not do so all the time. The word 'normal' has two distinct implications. The first is that it indicates adherence to some standard; the second is that this standard will in many cases be widely adhered to. It is owing to this second implication that Russell is able to treat 'normal' as a synonym for 'usual'. Yet he completely overlooks the first implication. But if the concept of normality has application to our judgements, then it indicates the existence of exactly what we need to find, a standard of correct judgement which will be very widely instantiated in the mass of judgements that people actually make. Now if there are normal ways of judging colours, the difficulty posed by Russell's argument vanishes. The way we normally judge things to be will of course correspond to the standard required for objective judgement, for that is what is implied by calling that way of judging normal. That way of

judging will also be a very common way in which to judge, although it need not correspond to the way people invariably judge, and need not correspond to the majority view in any particular case.

It is, in a way, absurd to doubt whether there is a normal way of making empirical judgements. For we have every reason to believe, from our ordinary experience, that in practice when a number of people look at some object they will normally (i.e. usually and rightly) agree about what properties it has, and they will normally (usually and rightly) all judge as they do for the same reasons, those reasons being the conditions for correct judgement which I have discussed and illustrated in this and the previous chapter. An underlying theme in this chapter has been that in practice we do recognise the existence of these standards of judgement, and that we do generally adhere to them. It is, however, another matter altogether to demonstrate that the standards we do all adhere to in practice are the standards which deserve to count as the correct ones. But now at least we know what we need to explain in order to show the connection between objectivity and the judgements we usually make.

One aim of this chapter, then, was to bring us to this point of recognising the significance of the notion of a normal judgement as combining the theoretical requirements for judgements to be objective with our common practice of judging as we do. But another strand of thought has also been running through this chapter. This was the idea that the judgements we make about our experiences are either judgements about the contents of a public world or else parasitic upon our capacity to make such judgements. In other words, our judgements are all necessarily ones which others could make and are thus all capable of being objective. This was suggested by noting how difficult it is for a description of an experience of colour couched in terms of neutral sense-data to be distinct from the public language we use to describe the public world and yet to be connected to that language, as it must be if it is to serve as the grounds for knowledge which others can also have.

Both these strands of thought stand opposed to any

epistemological theory which takes the individual's own experience, or his own mental awareness, as the foundation of knowledge and the basis for his claim to have made any judgement correctly. Although my arguments have not been aimed exclusively at the theories of the classical empiricists and their heirs, it is, I think, owing to the empiricist cast of mind in English philosophy more than to anything else that the concept of objectivity has been so much misunderstood, because the routes to knowledge and truth have been so much misunderstood. Therefore, the most direct way to reach a satisfactory account of the basis of objectivity in normal judgement, and the reasons for the logical priority of objective judgement over the subject's own experiences, is to turn to a traditional empiricist account of judgements like the ones I have been discussing in this chapter, in order to see where its weaknesses lie. This I propose to do in the next chapter.

CHAPTER FOUR

OBJECTIVITY AND EMPIRICISM: LOCKE

So far I have been arguing that we can make colour judgements (which I take to be typical and representative of judgements made on the evidence of the senses) only if we can make them objectively. And I have claimed that objectivity can be expressed in terms of how it is normal for us to judge, since this notion embraces both the analyses of objectivity we have made so far, the theoretical analysis in terms of justified standards of judgement and the practical analysis in terms of the standards exemplified by the way in which the majority of us do as a rule judge things to be. The tradition of English empiricist philosophy from Locke to Ayer has, however, taken a different view about the status of judgements made on the evidence of the senses, and in the previous chapter I argued against one aspect of it. As I suggested there, I believe the whole empiricist approach to the status of judgements to be mistaken and misleading, and in this chapter I shall try to show why this is. These arguments against the empiricist outlook will lead us directly to an opposing account of our capacity to judge objectively which will be far more satisfactory.

In this chapter I shall concentrate on Locke, because of his historical importance as a philosopher, but also because he was a great philosopher. If Descartes can be said to have been the originator of the sceptical method in philosophy, Locke was the first major philosopher to attempt to apply that method thoroughly to knowledge which is attainable through experience, and the first to begin to realise that language itself was an important factor to consider in such a study. Thus it was Locke who established the direction empiricist thought about knowledge and judgement was to take; and whatever his faults, he is a worthy spokesman for that movement.

Locke's 'Essay Concerning Human Understanding' (from which all subsequent references come) begins with an attack upon the doctrine of innate ideas. His reasons for the attack seem to have included a healthy distrust of the many metaphysical theories which depended on that doctrine, plus the conviction that a commonsense philosophy could show that there was no need for innate ideas anyway, since all knowledge could be shown to derive from experience. Locke seems to have thought that this latter point in itself constituted a disproof of the doctrine of innate ideas. He was wrong in this, since even if all knowledge were obtainable from experience, that would not preclude the possibility of one man's being born already knowing what other men only find out through the operation of their senses and their minds. I shall come back to this point later. The thing to notice here is that for various reasons, not all of them bad by any means, Locke thereby cut himself off from one type of consideration which might have rescued his epistemology from many of its difficulties. As it was, the attack on innate ideas was immensely influential. Isaiah Berlin has called it "historically, if not the first, the greatest blow struck for empiricism and against the vast metaphysical constructions which rested on axioms for which no evidence could be discovered."* However, as I shall argue, it also deprived empiricism of one possible way out of the blind alley of the quest for an account of knowledge based on experience alone.

Locke asks where the mind obtains all the materials of reason and knowledge, and answers "from experience" (II, 1,2). It is interesting that he considers reason as well as knowledge to be a product of experience. This may simply be an example of Locke's incautious phraseology, but there are other hints in the text which point in the same direction. The soul, says Locke, "comes by exercise to improve its faculty of thinking in the several parts of it." (II,1,20). Here we have the notion of a capacity being developed by use; but whether the capacity can be developed in different ways as well as to different extents

* 'The Age of Enlightenment' (Mentor 1956) p.40

according to the kind of exercise it has, Locke does not consider. I think he just assumes it cannot. But if the soul's capacity to reason cannot develop in different ways in different persons, then Locke is tacitly assuming that the faculty of reason has a determined structure, although different individuals may come to have a greater or lesser ability to use that structure according to how much exercise they give their minds, and a greater or lesser knowledge of it according to how carefully they attend to the operations of their minds, i.e. how attentively they reflect, and also, of course, on the stock of ideas of sensation they happen to have to reflect on. This is the first indication that Locke's epistemology begs an important question about the structure of our thought.

The unit of experience for Locke is the idea, though he sometimes uses this term to refer to types of experience rather than individual elements of experience. 'Idea' is defined by Locke as "whatever it is that the mind can be employed about in thinking."(I,1,8). Although the confusions in his use of this term are notorious, what Locke seems to have had in mind was that ideas all have a common function of being signs representing to the mind its objects, of whatever variety they may be. Unfortunately, this allows him to use 'idea' to refer to the content of immediate sense experience — what is given in sensation — and also to refer to what is perceived, that is, the recognition of ideas in the first sense as being of the things they are of. Thus, the various immediate sensory qualities of brown, hard, square, etc. are called ideas, but so is the recognition of these qualities as being the properties composing a table. This is another dangerous confusion, because it makes it easy for Locke to move from talk about sensations to talk about their objects without noticing that he has made any transition.

Ideas are either of sensation or reflection. The origin of all ideas is ultimately in sensation, for reflection is "only that notice which the mind takes of its own operations and the manner of them" when the mind is "employed about the ideas it has got". (II,1,4). This employment seems to be confined to the processes of com-

paring, reconstituting, assembling, abstracting and remembering, plus the incidental generation of some consequential passions. So ideas, once 'in' the mind, are rather like filing cards. They can be sorted and grouped in various ways, and some of the information on them can be collated on new cards, but the information on them cannot be supplemented or altered — though cards can, perhaps, be lost! So although the awareness in reflection that the mind has produced some new configuration of its materials constitutes a new idea, it follows that the original source of knowledge must be found in ideas of sensation. Locke is not always faithful to this account of the powers of the mind. Sometimes he recognises that it is more active than this account would indicate. But if the mind is more active than Locke has so far suggested, then my case against empiricism is strengthened. For the more the mind has to do to convert experience into knowledge, the less convincing is the claim that knowledge could be just a matter of absorbing experience as a lone individual. The claim is even less convincing if the possibility of shared knowledge requires different minds to have a common conceptual grasp on experience.

Locke describes ideas as being either simple or complex, and the latter are compounded in various ways from the former. He gives different accounts of what constitutes a simple idea. Sometimes he treats simple ideas as whatever cannot be analysed into anything simpler. Therefore any idea which is "not distinguishable into different ideas" is a simple idea (II,2,1). Simple ideas of sensation would therefore be by definition the basic units of experience. Sometimes, however, Locke employs a quite different criterion for the simplicity of an idea, and that is one in whose reception the mind is merely passive. (II,1,25. Also in II,12,1 Locke speaks of "those ideas, in the reception whereof the mind is only passive, which are those simple ones received from sensation and reflection". In the fourth edition of the 'Essay', Locke here inserts a reference to the mind being "wholly passive" in its reception of simple ideas.) Locke's two main definitions are not incompatible; for given Locke's

belief that the mind really does no more than sort and compare the basic units of experience, that it could not alter any basic unit of experience would follow analytically from his first definition of simple ideas, so the mind would have to be passive in respect of them.

This, then, is the equipment with which Locke thinks our knowledge of the external world can be constructed. It is not a very extensive kit, but it does not differ so very greatly from that which subsequent empiricists have employed to do the same job. Therefore objections to Locke's account are, broadly speaking, objections to the empiricist approach as a whole. There are from our point of view two major objections to a theory of this sort. First we may ask whether an individual, equipped as Locke describes him, could be aware of a world of external objects at all. Second, we may ask whether this world could be in any sense a public one. These two questions are really one and the same. For the definition of 'external object' must include the requirement that an external object be accessible to the senses of everyone, and not just something that figures in the experience of one individual, so an external object must be a public object. If a man insists that there is a tree in the middle of the field, but nobody else can see it, then whatever the man is aware of, it is not an external object. But though the conditions for a world of external objects and for a public world are the same, it is convenient to treat them as separate questions for the purposes of exposition.

On the face of it, there seems to be no possibility that the individual, equipped as Locke has equipped him, could distinguish between what is the case in the external world and what merely appears to him to be the case. (A fortiori, he cannot make the further distinction between how things really are and public appearances, which I introduced in Chapter Three.) For since he can never know of anything apart from his own ideas, he can never know whether they have any bearing on what the world is really like. Locke attempts to resolve this difficulty through his account of primary and secondary qualities. This account has been variously interpreted. In saying that

primary qualities resemble their objects, Locke has often been taken to mean that an idea of a primary quality gives us accurate information about an object in the world in respect of the quality in question. Berkeley certainly thought that was what Locke meant. Reginald Jackson, however, proposed a different interpretation, arguing that by primary qualities Locke meant qualities of objects in general, and by secondary qualities he meant the power which the primary qualities of objects have to produce sensations in us.* He suggests that if this is what Locke meant, at least most of the time, then primary qualities are, in fact, not knowable through sensation at all, but only by their powers to produce sensation in us — that is, by their secondary qualities. And these, being relational, do not resemble their objects, as Locke admits. Primary qualities would thus be postulates in a hypothesis (borrowed from Boyle's empirical hypothesis about the corpuscular nature of matter) about the causes of our sensations. On this interpretation Locke would in fact be saying that anything we could know about primary qualities would be a matter of inference from a theory, and not a matter of immediate awareness at all. If this was what Locke meant, then he came close at this point to appreciating the cardinal error of classical empiricism, which is to think that the individual mind can attain knowledge by merely assembling and comparing its own immediate experiences. But Locke seems to have thought that the threat this posed to other aspects of his own theory could be remedied by means of a priori considerations. He believed he could specify certain qualities which were also properties of matter, being "utterly inseparable from the body in what state soever it be"(II,8,9). These would serve as directly knowable primary qualities, without the need to resort to empirical hypotheses. Here, though, Locke falls foul of Berkeley's objection that he confuses determinate (i.e. particular) qualities with determinable (i.e. general and therefore definitional) qualities. This seems to have been a genuine blunder on Locke's part, and prevented him from seeing the correct conclusion to this line of reasoning, which is that all we can know of objects in the external

* 'Locke's Distinction between Primary and Secondary Qualities' (Mind 1929)

world solely on the basis of the sensory evidence available to the individual percipient is what must be true by virtue of the definition of 'object in the external world'; so, in effect we can learn nothing about particular objects in the world through experience alone, if experience is understood to be what Locke thought it was.

If, however, like most empiricists, we choose not to accept Jackson's interpretation of Locke's position, the situation is scarcely better. Then Locke's doctrine will be taken to be as follows: Some qualities in our experience do not resemble the objects in the world that produce them, but some do. The difference is that some qualities are recognisably relational; that is, they depend for their quality partly at least on the condition of the person who has the experience. Thus the experience of colour I have depends partly on the state of my eyes, the tastes things have vary according to the peculiarities of my palate, and so on. Such experiences cannot therefore be expected to give us correct information about what the external world is really like. However, there are some qualities which are not relational, and cannot vary according to the state of the percipient. Locke lists these qualities as solidity, extension, shape, motion or rest, and number. These qualities, being invariable, must, he thinks, be real properties independent of the observer's perception of them.

There is no 'must' about it, however. From the fact of the invariability of a quality we cannot draw the conclusion that it is not dependent on an observer for its particular invariable characteristic — and it must be remembered that we are still talking about the experience of a single observer. That observer might, for example, just happen always to see objects as extended in a particular way. It does not follow that the objects really are extended. This argument applies, <u>mutatis mutandis</u>, to any and all of the qualities Locke has listed. Furthermore, Berkeley's objection is again pertinent. We must note that the property of invariability only holds for determinable qualities. It would be trivial to assert that our ideas of primary qualities resemble real properties of things and that, as Locke puts it, "their patterns do really exist in

the bodies themselves" (II,8,17) if all that is meant by it is that every body must have some extension, number and so forth. But in that case, all that statements about determinable primary qualities could contain are assertions to the effect that particular bodies 'really' possess all those qualities in virtue of which they are considered to be bodies. Thus we arrive once more at the realisation that, on Locke's theory, knowledge of real objects, such as it is, turns out to be analytic and therefore uninformative. Locke does in fact assert at one point that "the particular bulk, number, figure and motion of the parts of fire or snow are really in them" (II,8,17), but admits a little later that he has gone rather further than he intended into natural philosophy. He recognises that here (as elsewhere) he is shoring up his philosophical argument with part of an empirical hypothesis.

On either interpretation of the doctrine of primary and secondary qualities, then, we can see that there is no case for saying that ideas as Locke conceived them contain informative knowledge of the external world. And indeed this is inevitable. For the basis of Locke's empiricism is that the individual's ideas not only form the foundations of his knowledge, but indeed they provide everything that body of knowledge can contain. Now, as we saw in Chapter Three, the terms 'internal' and 'external' mark a difference for us, but it is not easy to see how they could do so for the Robinson Crusoe figure, alone from birth, since all he has to go on are the qualitative differences between the different sensations he has.* But it is equally unclear why Crusoe should want to make that distinction, if it ever occurred to him to do so. For our notion of an external world is one in which objects are as they are independent of one's beliefs about them. Since there is no scope for that distinction to be made by Crusoe, there is no scope for a distinction between an external world and his internal impressions of it either. There is no basis for making that distinction in Locke's theory either, when all the individual can be aware of is the ideas within his own mind. So Locke's theory not only fails to provide us with information

* see pp.80-82

about the external world itself. It fails to provide us with the possibility of even making the distinction which would allow us to discriminate between external objects and the ideas in the mind of the individual.

The second difficulty I mentioned earlier was the question of whether the world as understood by Lockean man could be in any sense a public world. This amounts to asking whether the possibility that his experiences might be shared is one that makes sense. Experiences are always had by individual persons, of course, so it cannot be the case that two people should have qualitatively and numerically the same experience. But that is never what we mean anyway by speaking of two people having the same experience. There is a strong sense in which we can say that two people have (qualitatively) the same experience, because they can both have experience of the same thing; and there is an attenuated sense in which we can say that people have the same experience where the experience referred to is not of an individual thing but of a type of thing. So I can (in the stronger sense) see the very same table that you see, and I can have the same (sort of) pain that you had last week (for you described it to me). There are obviously important differences between what is meant by 'the same table' and 'the same pain', differences which turn largely on questions of the way in which criteria of identity are available for the objects of these different types of experience. These differences are not the issue here, for they present a problem of equal difficulty for Locke's theory in each case. What matters here for us is the way in which language and significant behaviour make it possible to recognise that two people are having the same experience, or else similar experiences. Obviously, my ability to describe my experiences is an important way in which you can find out whether my experience is at all like yours. But the significance of language goes deeper than that. For language (widely defined as any systematic body of meaningful signs) must be the vehicle through which the possibility of shared experience is grasped, since it is only language which has the capacity to convey the sense of experience. It is not only that I cannot know what another person is

experiencing unless I know what his words and actions mean. The very idea of there being such experiences available to the mind of another cannot occur to me unless there is a medium through which another's experiences can find outward expression. Words are central to this activity since, as Locke was among the first to notice, this representative or signifying function is one of their essential characteristics (although of course words perform many other functions as well — and even the representative function of language need not be anything like the sign-designating-idea model that Locke provides, which is really a supposedly interpersonal form of the idea-designating-raw-experience model which is how Locke sometimes describes sensation and reflection when, from time to time, he talks as if he is not identifying experience itself <u>as</u> ideas.)

Some reference to this function of language, then, becomes an essential part of any epistemology which purports to explain the possibility of there being shared experience. This is one reason why Locke goes on to discuss the nature of language; and for the same reason, the adequacy of his account is central to the question of whether his epistemology can describe a world of shareable experience, and hence a world in which objective judgement is possible.

For Locke, experience consists of the individual's stock of ideas. Words, then, must refer to these. He says that words are "sounds as signs of internal conceptions" (III,1,1) and "in their primary or immediate signification, stand for nothing but the ideas in the mind of him that uses them...nor can anyone apply them as marks, immediately, to anything else but the ideas that he himself hath" (III, 2,2). Hence the only immediate way of knowing what a word means is to have the appropriate idea available to one's mind and all mediate ways of knowing must finally come down to this. Now it is indeed the idea which constitutes not only the reference of the word, but its <u>sense</u> as well. But if this is the case, how can words be used to communicate at all? For in order to communicate ideas to one another, men must "suppose their words to be marks of the ideas in the minds also of other men" (III,2,4). And what grounds could there be for supposing <u>this</u>? A prerequisite

of the words being real words, that is, signs with meanings rather than mere noises that people make, is that men who communicate by using them must already have the same ideas in their minds, and must also annex the same words to these ideas. But there can be no way of knowing that the former is the case or that the latter is happening. For there is no way of knowing, except _via_ language, that men have the same ideas; nor, in these circumstances, would the fact that one man uses a word be a guarantee that he annexes the same idea to it that I do when I use that word. (These two points taken together also ensure that any disparities need not show up as inconsistencies in what we say.) So this part of Locke's theory collapses into circularity. And thus the question of what words an individual has associated with which (Lockean) ideas is undecideable, if indeed it is intelligible at all.

Locke evades this problem by slipping into his familiar empirical stance. There _is_, he believes, really a world of external objects which cause men to have the ideas they do have, and like causes produce like effects. Thus all men will have, broadly speaking, the same stock of ideas, occasioned by the passive reception by their minds of sensations emanating from the same objects by virtue of the secondary qualities or powers of those objects. Words will become associated with these ideas as they occur in different individuals by being taught in circumstances in which the same idea will be excited in the minds of teacher and pupil alike by the action of the same object which is present to them both. But of course the existence of such public objects and their ability to arouse the same ideas in different minds is the very point at issue, the very reason why language is being discussed in the first place. Locke's theory of language is not saved by recourse to an empirical hypothesis. And since he has failed to explain how there can be a shared experience, he has failed to explain how there can be a public world, because the latter is impossible without the former. And without the one or the other, there can be no distinction between subjective judgements and objective ones.

But although Locke's theory fails, we can learn from

its errors. For Locke assumes throughout that the meaning of a word, and also that to which the word refers in its primary signification, is located in the inner, and to others inaccessible, idea that the individual has. The idea is thus not only the referent of the word, but its relation to the word is what gives the word sense. Taking these points together, we can see that it is impossible for a strict empiricist account of a public experience to be given, precisely because no sign which could convey an understanding of what an experience is like from one person to another, could have its sense located in what is inaccessible to others — the individual's own stock of ideas — and still count as a sign for others. The sense of signs, which makes it possible for us to use them to refer, must be found in some area which is publicly accessible if communication is to be possible at all. But if private ideas are inaccessible, and public objects are only knowable as a consequence of the ability of individuals to compare experiences, then, paradoxical as it might seem, it cannot be the Lockean idea which is the determinant of what is to be communicated if experiences are to be publicly accessible. (Nor can it be the external object, as previous arguments have shown that Lockean man does not have a clear use for the word 'external'. Here we can see that in addition, an external object, being a public one, can only be known as a consequence of communication being possible.) So in some way it must be the mode of access, language itself, which determines through the sense its terms contain, what intelligibility shall be given to what is experienced.

I spoke earlier of words acting as signs for something else. It is natural to think of this relationship in terms of the sign being something added to a pre-existing referent; and certainly that is how Locke thought of the function of words. The tendency to think in this way is enhanced by the fact that the choice of the sign is arbitrary whereas the experience it is supposed to be attached to is thought of as already existing and, being a given which the mind passively accepts, already determined. However, if we step outside Locke's model of language for a moment, and

imagine the sign as the determinant of the otherwise inaccessible and so indeterminable quality of experience as it manifests itself in people in general (and so in oneself, since one is a person amongst others) we can see that the reverse of Locke's position is really the more convincing possibility. But from the standpoint of Locke's theory, with the idea determining the sense of words, communication through words is impossible; and with it the possibility of knowing that others are making the same judgement disappears as well. Yet such knowledge about the judgement of others is clearly possible. We take it for granted that it is possible, and we recognise it to be a requirement, in the typical case of colour judgements, for instance, for making the distinctions between appearing and really being. Since we cannot believe that our capacity to make such distinctions is illusory, the conception of the idea as being something which can be understood independent of the (public) sign we have for it, which lies at the very heart of Locke's epistemology, must be abandoned. (This might be taken as implying that sign and idea must really be one and the same, and that — to put a behaviouristic interpretation on that — there are <u>really</u> only signs. But I do not mean to imply that, and will in due course argue for a different interpretation. I do not mean that an experience which has no sign associated with it could not in some sense be said to <u>occur</u>, but only that it would not be intelligible as an experience. <u>A fortiori</u>, no experience could be intelligible as a shared experience if there are no words in which the experience could be described. The unintelligibility of experiences without corresponding signs will affect those cases too which I referred to earlier, where the experience in question is not one of something which can be shared, but only one where others can experience something similar. That I have a toothache, for instance, is an intelligible notion only if 'toothache' refers to something others can also have which is qualitatively like what I am having.)

I suggested at the start of this chapter that Locke's rejection of the possibility of innate ideas is a serious constraint on the feasibility of his theory. For there

are perhaps two explanations of the possibility of different minds coming to contain the same ideas which could be given without going too far from the terms in which Locke's theory is couched, although the theory would then be by no means what Locke intended, and it would no longer be an empiricist theory either. The first of these explanations is that independent objects might exist as the causes of ideas, no matter whose ideas they are. We have seen that Locke did make intermittent use of this view; but we have also seen that it is a scientific hypothesis rather than a philosophical theory. It cannot explain what gives rise to the belief that the same ideas might exist in my mind and in yours. It can only be introduced once we believe this, in order to account for our belief. As we have seen, there is nothing within the framework of the individual's stock of ideas, as Locke conceived it to be, which would even lead us to suppose that others might possess the same ideas as ourselves (and, it might be added, the conception of there being others like ourselves is similarly circumscribed). This explanation, therefore, can be abandoned from the outset.

The alternative is to incorporate the capacity of different minds to contain the same idea into a feature of the structure of thought itself. Yet Locke's approach began by ruling this possibility out. He believes that there can be no determined structure to thought itself, since if there were, it would manifest itself as a series of innate ideas. In Book I of the 'Essay' he argues that there exists no proposition either of reason or morals that absolutely everyone knows, which is doubtless true if by knowing we mean, as Locke does, having the ability to state the proposition known. He also claims, rather less convincingly, that no principle of reason or morals could be dispositionally innate; that is, no principle could be implicitly known, but the knowledge not be expressed or assented to until men come to the use of reason. For, he argues, it is not only difficult to see what could be meant by the claim that a person knows something he cannot state or assent to, but knowing a principle in this way is indistinguishable from coming to know it by experience, and much better explained by the latter suggestion.

In fact, Locke fails here to distinguish the ability to state or assent to a verbalised form of a principle of reason from the ability to use or deploy it. Not altogether surprisingly, Locke actually refers to this dispositional ability of the mind in the act of denying its existence. He writes:

"If by knowing and assenting to (innate ideas) 'when we come to the use of reason', be meant, that as soon as children come to the use of reason, they come also to know and assent to these maxims; this is also false and frivolous. How many instances of the use of reason may we observe in children, a long time before they have any knowledge of this maxim..." (I,2,12). Elsewhere Locke appeals to self-evidence or "intuitive knowledge" which he needs to explain the various abilities to operate which he ascribes to the understanding, such as the ability to recognise the similarities and differences between simple ideas. He admits that "It is on this intuition that depends all the whole certainty and evidence of all our knowledge" (IV,2,1). So if Locke is at times unwillingly prepared to admit that there are structural features of human thought which exist prior to experience and are not learned from it, his theory with the addition of some dispositionally innate principles is now more coherent at the expense of no longer being an empiricist theory in the strictest terms, for it no longer depends on experience alone as the source of both form and content of all knowledge. Furthermore, there would then seem to be no bar to introducing a number of other innate principles of reason in order to solve other problems within the theory, inclusive of the one we are interested in. I referred earlier to the confusion inherent in Locke's use of the term 'idea' to cover both what is given in sensation and what is perceived, since the latter but not the former seems to imply the existence of some thing which is the object of perception. Professor Mackie, discussing this, considers the possibility of there being an innate disposition to see things realistically; that is, he suggests that there might be an innate disposition to assemble concatenations of sensations as substantially based and causally related.* This would certainly solve

* J.L.Mackie 'Problems from Locke' (Oxford 1976) ch.7

the problems about causation and substance which Locke's epistemology also suffers from, but again at the expense of abandoning the principle that the only basis of knowledge is experience. To make a move of this sort is indeed to lean towards an important aspect of Kant's eventual response to empiricism. For to speak of dispositionally innate ideas is one rather loose way of referring to the possibility of synthetic a priori judgements.

From the standpoint of our concern with objectivity, however, even that modification to the empiricist approach to knowledge would fail to rescue Locke's theory entirely from its difficulties. For let us grant for the moment that there might be an innate disposition for experiences of whatever sort to be objectified as they are had. Now if this is to solve the problem of how different men can be said to have the same experiences, it can do so only by presupposing a form of sensibility, that is, a structuring of experience, which is the same for all men. But in that case, it is hard to see how we can avoid the conclusion that all men must judge alike in the same circumstances; for the same ideas will be excited in their minds as a result of their common sensibility. Consequently, the question will not now be how men can judge objectively, but how it is that subjective judgements ever come to be made. If a common sensibility determines that all men who observe a scarlet post-box have literally the same idea of scarlet as a result, then we can only explain the fact that someone might see the postbox differently from the others (as grey, for example) by saying that he and he alone had a completely different idea from all his fellows. (If he then puts his judgement into words, we can perhaps explain his error by saying that he has annexed the wrong word to the idea he has. But then how will we know that it is he who has made the error, and not the rest of us? We are then no nearer an explanation than we were before.) But the real problem is still to explain how his completely different idea originated, and how it can be reconciled with the putative existence of a common form of sensibility. For there are cases of a person seeing something different from others, and yet proving by consistent usage elsewhere that he does

understand the terms he is using, because in other cases he does annex the word to the idea as we all do. We touch here on a particular problem for any theory which tries to impose a priori mental structures on our capacity to make empirical judgements, for they seem to leave no room for subjectivity in particular cases, and leave the problem of error exactly where it was before. This is a major and, I think, unresolved difficulty in Kant's solution to the problem of knowledge which empiricism bequeathed to him. Thus, at least in the simple form in which it has been presented here, the introduction of dispositionally innate ideas solves one problem only to raise another. I have been unable to find or think of a form this argument might take which would overcome this difficulty; but in view of the fact that a more satisfactory solution to this problem will be given in the next chapter, I do not think I am guilty of an important failure here.

An alternative response to the difficulties in Locke's account of the possibility of public agreement in judgements (and one which may have seemed to be implied at one or two points in this chapter so far) is to reject his account of inner experience altogether as irrelevant. If the possibility of public agreement in judgements makes us despair of explaining how the sense of words can come from the individual's own, and to others inaccessible ideas, then it seems that the sense of the words must come from something which is itself public. That something will not be the relation of the word to its reference, the object itself, or its relevant qualities, since the possibility of public agreement in judgements is a precondition of there being public objects or public qualities to refer to at all. Therefore (since that is all that remains) the sense of words used in public judgements must derive from the way in which people actually use them when they make these judgements. Not any usage will do, though. Usage has actually to be correct, or our words might mean anything or nothing. But providing that we can establish what counts as correct interpersonal judgements in relation to any particular concept (such as redness) so that we know when it is correct to say 'that is red' (on this theory, when

people in general do say it), we have established the sense of the word 'red', and without the need to refer to ideas at all. The individual's idea, whatever it is or is like, drops out of consideration as an irrelevance.

Part of the trouble with this theory, however, is the requirement that the judgements people make should actually be correct. For we know that what most people say does not in itself establish their correctness. If agreement in judgements is the only basis available for the sense of words, we are entitled to ask whether this agreement is justified agreement, or whether it is merely a consensus view. This brings us back again to the problem we faced at the end of Chapter Three, but now with an additional difficulty. For not only do we now lack the idea of there being a justification for our judgements which we could regard as sufficient, so that we would be reduced to asserting that judgements were objective on the inadequate basis of how it is usual to judge, rather than on how it is normal to judge; but in addition we have now located the sense of the terms we employ in our judgements in the agreement itself. Aberrant, mistaken or subjective judgements would then be entirely incomprehensible. It would not merely be the case that someone who said 'this is yellow', about something which it was generally agreed was red, would be uttering a falsehood. He would be uttering something which could not be intelligible, either to anyone else or to himself, since the criteria for correctness and intelligibility within this theory are one and the same. So the price of accepting this alternative to the empiricist approach would be even greater than the other alternatives we have considered. Other theories have made disagreement in judgements puzzling; this theory makes them impossible. For genuine disagreements require there to be meaningful alternatives, and this theory allows for none. Perhaps this theory might be sophisticated to the extent of allowing some sort of distinction between correctness and intelligibility to be made; but any attempt to do this could only be effective to the extent that it contains some grounds other than the agreement of the majority for coming to make a judgement. And once we admit other grounds for judgement, or even for understanding

the concepts employed in making judgements, we are inevitably driven to take account of the individual's experience again, which really negates the whole point of having the theory in this form. The mistake of empiricism is to try to make the individual's experience the sole basis for public judgement. The mistake of the radical alternative we are now discussing is to rule out the role of individual experience entirely. In seeing this, we see also that an adequate theory of objectivity needs to explain not one but two things about public agreement in judgements. It needs to explain not only how judgements can count as actually correct, but also how the concepts employed in making the judgements come to have application, which amounts to explaining how the terms we employ in judgements come to have sense. Locke wrongly located the answer to both questions in the relation between the judgement and the idea in the mind of the individual subject of experience. The radical alternative tries to locate the answer to both questions in the relation between the individual's judgements and the judgements of us all. Both accounts are defective because neither gives full weight to both these aspects of judgement via the senses. If subjective judgement is to be allowable, the individual's actual experience must play some part in his judgement. And if objective judgement is to be possible, and judgement as a whole is to be meaningful, interpersonal standards of judgement must impose themselves in some positive way on the individual's judgements at the same time.

But though the new question we have raised about how terms used in judgements have sense, seems only to be a further complication thrown up by consideration of the failure of empiricism to explain how objective judgement is possible, it really provides a key to the solution we are seeking, together with the notion of normal judgements which we arrived at in the previous chapter. For when we come to see just how words for our experiences get their sense and are able to function in a public domain, we shall also see why it is normal to judge in some ways rather than others. In other words, the reason why we are justified in judging things to be as most people, most of the time, do judge

them to be, will emerge from an adequate account of how the terms employed in our judgements get their meaning.

This, however, is a topic for the next chapter. The aim of this one has been to locate the inadequacies of Locke's account of judgement with respect to the problem of objectivity, and so to lead us to the point where we can consider a theory which will overcome those inadequacies. In the process, we shall also overcome the difficulties inherent in the contrasting behaviouristic theory of judgement which I have sketched in over the last few pages.

CHAPTER FIVE

RULES AND EXPERIENCES: WITTGENSTEIN

The previous chapter showed us the need for the words for our sense-experiences to have the same meaning for each of us if objectivity is to be possible. At the same time we saw that the extreme, but obvious, form of this doctrine, the claim that the individual's having of inner experience is unconnected with his possession of the concepts which seem to be based on those experiences, involves major difficulties. For belief in the objectivity of judgements is degraded to faith in the rightness of a consensus of judgements, which is unsatisfactory because calling a judgement objective does involve more than people merely agreeing in their judgements. The agreement must be justified agreement, as we have seen. Furthermore, I have argued that the extreme form of this doctrine makes interpersonal judgement possible at the cost of rendering meaningful disagreements in judgements impossible. Thus, even if objectivity were possible according to the terms of this doctrine, it would be at the cost of making it impossible for there to be subjective judgements as well. So we still face considerable problems in explaining how objective and subjective judgements can occur in the way they do.

The work of the later Wittgenstein is highly relevant to these difficulties. It is widely recognised that in the 'Philosophical Investigations' he attacks the notion that possession of a concept depends on acquaintance with an inner, logically private, sensation or idea.* This has led some philosophers to interpret him as advocating something like the opposite view, described above, where concepts have nothing to do with inner experience.** But in spite of there being some evidence that Wittgenstein may sometimes have been tempted to think in that direction, I

* tr. G.E.M.Anscombe, 2nd. edition 1958 (Blackwell)

** see for example G.Pitcher 'The Philosophy of Wittgenstein' (Prentice-Hall 1964) p.298ff; J.J.C.Smart 'Philosophy and Scientific Realism' (Routledge 1963) esp. Ch.4; N.Malcolm 'Dreaming' (Routledge 1959) passim.

do not think that his argument against basing meaning on acquaintance with inner experience needs to be interpreted so radically. In what follows I shall argue for an interpretation of the Private Language Argument which does not require us to deny either the existence of inner experiences or that they have a role in the making of judgements. At the same time this account will dispel the difficulty raised in Chapter Three, where we encountered objections to our belief that how we ordinarily judge, as standard judges in standard conditions of observation, is no guarantee that our judgements are objective, because there needs to be justification for our belief that how we ordinarily judge is how it is normal to judge. Wittgenstein's work offers us such justification.

The Private Language Argument is deployed by Wittgenstein as part of a wider strategy intended to alter our way of looking at the concepts of meaning, understanding and communication. This intention, coupled with the writer's aphoristic style, makes exposition difficult, and it is often hard to focus on the application of his ideas to one particular issue. Although most of the following material is drawn from 'Philosophical Investigations' ss243-304, other references will be introduced from time to time where this is helpful.

First of all we must note that a private language, as Wittgenstein uses the term, is a <u>logically</u> private language, and not one that it just happens that no-one else understands. We are invited in ss243 to imagine a tribe of people who use language to speak only to themselves. This is not a true private language, since an explorer could come to understand it by observing how the actions of the members of the tribe fitted in with what they said. He could come to predict their actions from their words by this means. This is to assume, of course, that the explorer could make sense of what the tribesman did, but this is an assumption that Wittgenstein is prepared to make. That human behaviour in general makes sense is not for him an empirical assumption, but one by which the whole notion of 'making sense' is underpinned. He makes this assumption clear in ss206, where he says: "The common behaviour of

mankind is the system of reference by means of which we interpret an unknown language." The tribal language of ss243 is thus in essence a public, not a private language. But, Wittgenstein asks, can we also imagine a language being used only to signify inner experiences? "The individual words of the language are to refer to what can only be known to the person speaking; to his immediate private sensations. So another person cannot understand the language." Nor are the words in this language to be linked to any "natural expressions of sensation" which might enable another person to understand the words (ss256). The link between word and sensation is quite arbitrary. "I simply associate names with sensations..." Wittgenstein's purpose here is the negative one of making it clear that the private language cannot be taught to another or learnt from another. We are so used to the everyday correspondence between word and sensation that it is easy to overlook the fact that a private language must be invented for oneself or else just be with one from birth. Locke certainly failed to appreciate this point; and indeed, the language that Wittgenstein is describing here is virtually identical to that described by Locke in the 'Essay' as the model of our language, although Locke, as we saw, employed various shifts to get round the problems he created for himself in making words stand for ideas in the mind of the individual. One such, it will be remembered, was to suppose that the language could be learnt by making the (unjustifiable) assumption that the object itself could cause the same idea to be in the mind of pupil and teacher. I have already argued, in Chapter Four, that a language composed of such words as this could not be used to communicate with others; and Wittgenstein also makes this point briefly in ss257. So such a language is useless for the expression of judgements which others could either agree or disagree with. Wittgenstein's case is a still stronger one, however, for he argues that such 'words' could not form part of a language at all. To call these signs words is to misrepresent entirely what language is like.

The process of naming a logically private experience is now examined. Wittgenstein concentrates here (ss257 ff)

on the example of naming a pain, but any sensation could be substituted for pain as the example. Elsewhere he talks about sensations of colour in just the same manner, and his use of the word 'sensation' is, I think, meant to apply to any mental state or occurrence where the test of recognition or understanding seems to be a looking inwards at a sort of 'inner picture' which serves as an exemplar for the state or process. So he certainly means his argument to cover all those cases where we would talk of sense data, pure phenomenal experience or a Lockean idea of sensation. We must beware of letting the specific (and in some respects slightly misleading) example of pain blind us to the real point of the Private Language Argument, which is to deny that the private linguist can possess any concept which is understood solely by being associated with the logically private sensation which the private linguist has. Coupled with this denial is another, the denial of the psychological tendency to feel that not only is the private sensation the exemplar of the meaning of the term associated with the sensation, but that the understanding gained by this means is as complete as it is possible for understanding to be. Both ideas are contained in the rather cryptic remark which Wittgenstein puts into the mouth of the private linguist in ss264: ""Once you know what the word stands for, you understand it, you know its whole use."" There is also an echo of both ideas in ss274: "It is as if when I uttered the word I cast a sidelong glance at the private sensation, as it were in order to say to myself: I know all right what I mean by it."

In order to expose the mistake underlying both these ideas, Wittgenstein introduces the device of the diary (ss 258) used to record the recurrence of a particular (logically private) sensation. Every time the sensation S occurs the diarist will make a note "S" in his diary. (The diary is not, of course, a way round logical privacy. There is no intention, as there is no possibility, that anyone else should understand from the diary what the sign written in it signifies. The point of introducing the diary is that it serves to show up the process which is supposed to be going on when a private language is in use.) Wittgenstein

first points out that "S" cannot be defined. Its only sense lies in its relation to its reference, the sensation itself, since the sense of "S" is given by the act of making a reference to S. This is "a kind of ostensive definition" but not in the ordinary sense. Instead of pointing to the sensation, the diarist does what he supposes to be the mental equivalent of pointing — "I concentrate my attention on the sensation — and so, as it were, point to it inwardly". Thus the diarist seems to have created a record of S for future reference. But this is an idle procedure unless it has the result of allowing him to identify S when S occurs again. Wittgenstein remarks that ""I impress it on myself" can only mean: this process brings it about that I remember the connexion right in the future." In other words, the process of concentrating on S while uttering "S" must transform the sign which will mark not only that occurence of S but future ones as well, so that "S" becomes a word which is able to express a concept. Now it is fundamental to the possession of a concept that there are rules for the application of the word that expresses it. This is because a sign which is supposed to stand for a concept yet could be applied to anything in any circumstances would not indicate a concept at all. The user of such a sign would be showing, by his way of using it, that he did not in fact possess a concept here, that he did not understand the concept he believed himself to have by virtue of having a sign for it. Now the private diarist is not in a position to use the sign "S" rightly or wrongly at all. To justify his use of the sign he can only check it against his memory of what "S" is supposed to 'stand for'. The difficulty here has nothing to do with the fallibility of memory, as is sometimes supposed, for the diarist is in just as bad a situation if he merely believes that what he now feels or experiences is another occurrence of S. Saying that he merely believes it implies that this time he has got it wrong, or might have got it wrong. And the diarist's ability to make the connexion between S and "S" right or wrong is precisely what is in doubt. In ss265 Wittgenstein says "Justification consists in appealing to something independent." Now in some circumstances, an appeal from

one memory to another (or even concentrating on the memory one has) can amount to a justification for the correctness of the original memory; but this is only because memories are themselves checkable against the facts, so that we know that there is a difference between remembering right and remembering wrong. But in the diarist's case this sort of distinction cannot ever be made. For only his memory, and nothing else, can tell him whether he is following the rule for connecting S to "S" correctly or not. Even if the diarist retreats one stage further and claims that he has his own impression of what counts as 'following a rule' or 'getting a designation right', the situation is no better. For it is still the case that "Whatever is going to seem right...is right. And that only means that here we can't talk about 'right'." (ss258).

The private diarist, and thus the private language user, has the impression that private ostensive definition is possible because in fact he really does already have the concepts he is supposing himself to be defining. This is because he already speaks a real language, and has in fact gained the concepts, and the rules which make them meaningful, from that language. Ostensive definition in general only works when we already know something of the use of a word in a language, but are uncertain about how to apply it. Earlier in the 'Philosophical Investigations' Wittgenstein wrote: "Ostensive definition explains the use — the meaning — of the word when the overall role of the word in the language is clear." (ss30). So the technique works if, say, a person already speaks one language but needs to know what colour a certain word signifies in another language. For he has a grasp of the appropriate concepts of colour and colours through his mastery of his own language. The technique of ostensive definition also works when we teach a child its first language, not because we can guarantee that the child has the right (private) sensation before his mind when we utter the word (as Locke had to suppose) but because the child's future use of the word, and his future behaviour in connection with it, will show whether he has grasped the concept or not. If he has, his use of the word will, by and large, correspond with ours, for

reasons I shall shortly come to. However, here I must stress again one of the lessons of Chapter Four. The agreement of the child's future usage of a word with our own does not mean that the sense of the word for him, and of the concept that lies behind its proper use by him, is the same thing as his ability to use the word as we all do. He uses a word like 'red' on account of his experience of redness; but he needs to learn the public rules for using 'red' (and so grasp the concept involved) in order to make these experiences intelligible to himself and communicable to others. And that is what the ostensive process is teaching him. Wittgenstein is not ruling out inner experiences (or any sort of personal experience) in his attack on private language theory. He is ruling out one particular model of the relationship between the individual's sensation and the intelligible public sign for it. In doing this, he also shows how firmly the understanding of the individual's experience is rooted in the public domain, whether that understanding is of one's own experiences or of someone else's. This is an insight which is central to my argument.

According to Wittgenstein, then, the private linguist's conception of what it is to give a word meaning is such that the result of his activity could not amount to a language, even for the private linguist himself; a fortiori, it could not be a language which he could have in common with other speakers. This last point has indeed already been established in the previous chapter by other arguments. My reason for turning to Wittgenstein's arguments on the topic is that they enable us to see why this should be so. In rejecting the model in which words derive their meanings from private sensations which serve as exemplars, Wittgenstein emphasises the fact that it is an essential condition for a word's having meaning at all that its application to states of affairs should be rule-governed, that there should be a significant difference between using the word rightly and using it wrongly. He further stresses that the notion of a rule is only intelligible, and the distinction between right and wrong usage is only significant, where the idea of using a word correctly is a public

one, and hence where the concept that the word signifies is one that others also have. For, as we have seen, a private rule would be no more than an impression of a rule, which one can neither follow nor fail to follow. All this amounts to a radical shift of emphasis in our understanding of the theory of meaning. It forces us to abandon our conception of meaning as the relation between object and designation, even in that area, inner experience, where it seemed that no other explanation would possibly suffice. This is Wittgenstein's real motive here, and it explains why, when his protagonist asks: "Aren't you at bottom really saying that everything except human behaviour is a fiction?" Wittgenstein replies: "If I do speak of a fiction, then it is of a grammatical fiction." (ss307). Wittgenstein was insistant that we are often bewitched by linguistic ("grammatical") habits which we carry over from our ordinary ways of talking and then assume must be the form explanations will take when we do philosophy. The linguistic habit he is concerned with here, however, is not the fact that we refer to inner experiences as things that occur, but that we assume that all reference to what is public must gain its sense first from its relation to the individual's inner, private sensation — the error of empiricism. The question of the inner experience's existence or occurrence plays no part in the false theory's exposure which Wittgenstein has presented. It is only the role of inner experience in the theory that he rejects. That leaves open the question of what place, other than as an exemplar, one's own individual experiences occupy in a theory which can explain how public agreement in judgements can occur. I shall return to the question of subjective experience later in this chapter. Our main concern at the moment, though, is still objectivity, and the problem of why, when people judge as they ordinarily do, they are normally justified in the claims they make.

I have already indicated that Wittgenstein lays stress on the notion of following a rule in his account of meaning and condemns private languages because such languages cannot have rules. His point about rules, as it occurs in his main discussion of private languages (ss243-304), is often

the subject of misunderstandings, partly because of the way Wittgenstein uses the words 'right' and 'correct', and partly because his readers take insufficient account of the discussion about rules which immediately precedes this section of the book. Thus sometimes Wittgenstein is taken to be advocating the view that to use a word correctly amounts simply to using it as others do, because of his remark about remembering right in future, and particularly because he rejects the private diarist's exercise of checking one memory against another by saying "No; for this process has got to produce a memory which is actually correct." (ss265). This tends to convey the impression that the public rules for the application of a sign must result in a judgement which is correct in every case; and hence that so long as one follows the rules in judging (which is taken to mean 'so long as you judge as others do') one will simply judge correctly. Obviously, this is the same doctrine which I discussed at the end of Chapter Four, the belief that correctness in judgements simply amounts to judging as the majority do. We have seen that this view is inadequate; but it is also not at all what Wittgenstein is saying. His real point about the need for an independent check, which is a check on whether the public rules for the application of a word are being followed, is, I think, this: only the possibility of an appeal to "something independent" can allow us to talk of a check establishing that a word has been rightly or wrongly used. The private diarist cannot do this, of course, which is why his model of how words acquire meaning is seen to be absurd. But, on the one hand, it does not follow from this that my judgement can only count as a correct one if we proceed to check it against the judgements of others there and then, so that its correctness is assured by its being like their judgements (if, indeed, it is). And, on the other hand, the fact that a person uses a word just as others do may be indicative of his mastery of the concept involved, but is not identical to the understanding of the concept that this mastery implies.

Wittgenstein's idea here, then, is that there would be no room for the possibility of rightness or wrongness

in the judgements made in a private language, and not the point that either correct or meaningful judgements consists merely in judging as others do. The fact that one must as a rule judge as others do if one is to judge correctly or use the terms employed in judgement in a meaningful way, emerges from a quite different part of his argument; and when Wittgenstein says that an independent check has got to produce something that is "actually correct", he is assuming (often wrongly) that the reader has followed his thoughts from the discussion of rules in the sections before ss243, and in particular the very compressed thoughts in the bridging passage which occurs at ss241-2.

Section 241 begins with Wittgenstein's interlocuter suggesting the view Wittgenstein wants to guard against: "So you are saying that human agreement decides what is true and what is false?". This is emphatically denied. "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 form of life." Here Wittgenstein is stressing the important fact that a distinction must be drawn between the grounds for making particular judgements (what human beings say) and the way in which human beings come to understand the concepts which they deploy in making those judgements. This foreshadows the argument, of which the Private Language Argument forms a part, that it is one thing to judge (for example) that something is red, and to judge this on account of the experience one is having, and another thing altogether to come by the concept of redness and to be able to communicate by using words for such concepts. (This distinction is what allows us room for subjectivity in our judgements, and allows us sometimes to be wrong without being unintelligible, as I shall argue in some detail later.) The mastery of a concept through learning and following the public rules for the use of the word that expresses it, and the application of that concept by the individual to his own experiences, are shown to be necessarily different by the Private Language Argument. But though they are different, they are connected in an important way. It is left to ss242, with its rather cryptic analogy, to remind us that there is a connection as well as a difference, and to indicate how the connection functions.

The analogy concerns measurement, but we may take this as a representative example of any kind of judging. Wittgenstein writes: "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." The first sentence here makes the distinction between understanding a word, and actually judging that it applies here, which we have just noted. It is the second sentence which contains the crucial insight. If we did not all make measurements with (broadly) the same results as one another, he is saying, the concept of measurement would not exist for us. In order to appreciate the full significance of this idea, we must recall what Wittgenstein says in his Private Language Argument about rules. Having a rule is what enables us to distinguish proper applications of a word standing for a concept from improper applications of that word. This limitation of the possible applications of the word limits the concept and makes it meaningful, because it cannot then have application to anything and everything. But a rule is only a rule if it is something independent. I cannot decide for myself what the rule is to be, and decide too what counts for me as following the rule. This is why a private language is not a language, for the rules for the application of signs to sensations are just "impressions of rules" (ss259). Thus the 'signs' in that language do not, and cannot, have a sense. Now the judgements of each of us provide that independent check for others which any rule must have in order to be significant, and not just the individual's impression of a rule. Thus words can have meanings for us, and the concepts they stand for can be intelligible. But the independent check will only be an independent check, and the words in the public language will only have meaning for us all, and the concepts they stand for will only be intelligible, if the judgements that all of us make provide a constant standard. For only something constant could constitute a rule which it would be possible to follow or to deviate from. Anything less than what Wittgenstein calls "a certain constancy" would be too little to amount to a rule that anyone would be able to follow;

and without a rule to determine correctness or incorrectness, words and the concepts they stand for would lack any sense. So in essence, Wittgenstein is pointing out that words and the concepts they express require there to be interpersonal agreement in judgements if they are to have sense, although he is not saying that the sense of words or their concepts will determine the content of any particular judgement made by any particular individual on any particular occasion. This first thought is contained in ss242, when he says: "If language is to be a means of communication there must be agreement not only in definitions but also....in judgements." The connection between agreement and rules he has pointed out already in ss224, saying there that if you learn what the one means, you have also learnt the meaning of the other. And the second thought is contained in the denial that human agreement decides what is true and what is false, as we have already noted.

It may be wondered why Wittgenstein calls the agreement in judgements which gives sense to the concept of measurement "a certain constancy", and not simply 'constancy'. This is quite easily answered, however. It is not that Wittgenstein is expressing any hesitancy here. It is rather that he is alluding to the fact that although there must be agreement in judgements as a whole if expressions of the requisite type of judgement are to make sense, yet because individual judgements are made on the basis of the individual's experience (though his appreciation of the rules is required too in order for him to make sense of the content of his own experiences), the individual is not bound on every occasion to judge as others do. He can, of course, judge wrongly. This, however, cannot be the regular state of affairs. If the individual regularly judged differently from others, we should say that he did not understand what he was judging, or what counts as judging in this sort of case. And of course if people in general did not regularly agree in their judgements, the whole structure of intelligible concepts, and the very intelligibility of the language that expresses them, would collapse; indeed, if disagreement were chronic, the intelligibility would never have come about in the first place. We can see, then, that

when Wittgenstein refers to a certain constancy in the results of judgements, he is in fact indicating that notion which I said, at the end of Chapter Three, was what we needed to establish in order to give a proper account of the basis of objectivity. For what Wittgenstein is referring to here is the notion of there being a way in which it is normal to judge.

Here, then, at the beginning of the argument about why words in a private language could not have sense, is the key to the analysis of objectivity. Language only makes sense, and the having of concepts is only possible, if it is the case that people do normally all agree in their judgements. It is this normality of agreement which makes it possible to follow a rule, which in turn enables us to distinguish both right judgement from wrong judgement and sense from nonsense. Now we can see why the doubts raised at the end of Chapter Three were unfounded. These doubts centred on whether the way that people usually do judge could count as the standard which would justify their judgements, and so allow those judgements to count as objective. It is now apparent that this must be the case, even though judging aright does not simply amount to judging as others do. We could not say that people as a whole are wrong in their judgements without also saying that they lack understanding of what they are saying. And while this might be conceivable in a particular area of judgement, especially if the area is a recondite one, it could not be so of a wide range of human judgements if people are to be said to understand the language they speak. But at the same time we can see that it is always possible for the individual's judgement to fail to conform to the normal judgement on occasion, and still refer to something in such a way as to be intelligible to the person who makes the judgement and communicable to others. For the individual's judgement will have been made on the basis of his own experience; and providing he normally judges as others do, the concepts which make his experience intelligible and communicable are still his to command and deploy.

There is one other point about Wittgenstein's remarks here which deserves clarification. He speaks of the con-

stancy of results in measuring as what partly determines our understanding of the concept of measurement. Again, this may suggest that there is some further consideration lurking in the background, but it seems to me that this is no more than a piece of Wittgensteinian shorthand for a rather simple point. There are all sorts of different rules which have to be followed when a particular concept, such as that of measurement, is deployed in judgement. Most of these rules serve to differentiate one concept from another, or one sort of concept from others. One measures the length of a line, for example, but not the length of an area. And one measures the length of the line by placing a rule along it and reading off a figure, not by waving the rule in the air and making up a figure. What one does shows whether one understands how to measure. Most rules for the application of concepts, then, serve to differentiate one concept from another and specify ways in which that concept should be applied, and ways in which doing or saying something can fail to be a proper application of that concept, though perhaps it might be appropriate to some other concept. But the rule that Wittgenstein has singled out for mention is a rule of a different sort. Every concept whatsoever, if it is to be a concept, must be applied in such a way as to follow this one special rule. Every concept, whatever else counts as a requirement for its proper application, must be alike in this one respect. There must be normal agreement among people in the results of their judgements if the concept they are using is to be intelligible. This is Wittgenstein's point; but he is acknowledging the existence of the other rules which differentiate particular concepts in passing, as it were, by saying that the special rule is only part of what we have to understand in order to understand the particular human activity called measuring.

I have now outlined what I believe to be the only type of theory which can explain how the judgements that people normally make are objective, but which also allows for the possibility of there being meaningful judgements, grounded in the individual's own experience, which need not be in accord with the judgements of the majority. The only

alternatives to such a theory would be the two possibilities that I have already considered and rejected. For the fact that some interpersonal judgements but not others are justified could not be explained by basing justification on mere agreements in judgements; and that theory was proposed as the antidote to the empiricist's attempt to base the justification of judgements on the hopelessly inadequate ground of the individual's awareness of his own sensations. It was Wittgenstein's peculiar achievement to bring these two opposing views into a coherent synthesis. His work has provoked a vast body of comment and criticism, but it would be as unnecessary as it would be daunting to turn our attention to it here. Much of it has been, as mine in this chapter has been, an attempt to clarify and restate rather than to oppose. For a theory of objectivity broadly on the lines of the one I have tried to extract from the 'Philosophical Investigations' must be correct, inasmuch as any theory which can explain why everybody's judgements are normally, but not always, justified must do so by showing that the public's way of judging bestows intelligibility on the individual's experiences, and hence on the individual's judgements which he makes on account of having those experiences, yet at the same time <u>not</u> determining the experience or the judgement and so collapsing into a theory of justification by consensus of opinions. And that, as I have argued, is what Wittgenstein has been able to do, by maintaining the reality of the individual's experiences as a source of knowledge while making the intelligibility and communicability of those experiences something belonging to, because derived from, the public domain. It now remains for us to examine more closely the function of the individual's experience in the theory I am putting forward, so as to show more clearly how our experience can lead to our making subjective as well as objective judgements.

It will be remembered that towards the end of Chapter Four I claimed that theories purporting to base the notion of objectivity <u>solely</u> on interpersonal agreement in judgements were open to serious objections. Depending on the form the theory might take, it would exclude in one way or another the possibility of meaningful disagreement in judge-

ments between people, and hence deny that there could be meaningful subjective judgements at all. In spite of this difficulty, however, it has sometimes been suggested that inner experience plays no part in either the acquisition of concepts or their employment in judging. These two suggestions are different, although they have not always been clearly separated. Both views are counter-intuitive, of course, and both appear to rest on that misunderstanding of the Private Language Argument which I have tried to lay bare in this chapter. The first view, that inner experience need play no part in our coming to possess a sense-related concept, is the more radical thesis. For since coming to possess a concept is ipso facto to come to possess the capacity to make correct judgements of the relevant sort, the consequent ability to make judgements without reference to any inner experience is also implied. But hardly less extreme is the suggestion that although the having of an inner experience gives rise to the making of a judgement, so that in that sense the occurrence of judgement without inner experience is inconceivable, the inner experience cannot provide a reason for making the judgement. Therefore we have to look to interpersonal agreement in judgements as the sole form that all justification of judgements must take, with all the attendant difficulties for our being able to explain the occurrence of subjective judgements which I have mentioned. The first of these views has been put forward by Prof. J.J.C.Smart, and the latter has sometimes been attributed to Wittgenstein himself. I shall examine each view in turn.

In his 'Philosophy and Scientific Realism'* Smart sets out to present a theory of knowledge compatible with the findings and methods of science, and in this inner experience has no part to play. Taking colour judgements as representative of judgements involving secondary qualities, he attempts to describe a theory of colour in terms of discriminatory reactions, without any reference to inner experience at all. He believes it to be possible to form, and hence to employ, a notion of colour solely on the basis of the ability of standard judges to make discriminations. I shall not go

* Routledge 1963

into his account of standard judges, since it does not differ substantially from the one I have given in Chapters Two and Three, except that it is careful to make no reference to perceptual ability, referring instead to discriminatory ability. The significant feature of Smart's view is that he believes the ability of standard judges to make discriminatory responses is itself sufficient to allow the formation of concepts. Thus, to cite his example, a blind man could in certain circumstances come to have as proper and complete a concept of colour as a sighted person. In this connection Smart remarks "the idea that a congenitally blind man could not understand colour words is connected, I think, with a pre-Wittgensteinian view of meaning not as 'use' but as some sort of mental experience, which evokes and is evoked by a word" (op. cit. p.81). (This rather simplistic view of Wittgenstein's arguments certainly suggests that Smart takes a more radical view than I have of just what Wittgenstein is denying when he challenges the central role of the individual's sensations in our account of human knowledge!) In order to illustrate his view, Smart asks us to imagine a situation as follows: there is a race of congenitally blind people who have as slaves a race of people with ordinary vision, a large number of whom (perhaps all) are able to make standard judgements. However, the slaves are forbidden to use words for colours, even to each other, so they have no understanding of what such words might mean. Nevertheless, Smart argues, through their slaves the blind master-race could come to have an understanding of colour no different from what they would have if they possessed normal human vision. Naturally, the masters would not be able to make colour judgements by the same means as us; they could not identify the colours of things directly by vision alone. But they could make these judgements indirectly, using their slaves' discriminatory abilities but not, of course, their slaves' conceptual grasp of colours, since the slaves have none, or at least none that they can communicate to their masters because of the ban on speaking of colours. Smart aims only to show that concept-formation requires no reference to the having of inner experiences; but I think it is

clear that his argument would extend to judgements as well, providing only that the person making the judgement has means, direct or indirect, of making the appropriate discriminatory responses. (We might imagine the slaves in Smart's example being replaced by a range of mechanical sensors, for example.) The masters become aware of discriminatory responses to colours in the following way, in Smart's account, though obviously other ways would be equally possible. Quantities of variously dyed pieces of wool are provided by the masters, who get their slaves to sort the wool in 'some significant way'. The wool pieces are tagged or marked in some (non-visual) way known to the masters but unknown to their slaves. So by examining the sorted piles of wool, the masters will learn of the discriminatory categories that are represented by the different piles of wool, and can go on to give names to these categories. The masters can also learn of the existence of a colour continuum from the ease or difficulty with which the slaves can distinguish between wools of similar shades of colour, and by adding natural objects to the piles of wool can discover that these objects also display the features which distinguish the different sorts of dyed wool. Thus they can come to apply their colour-language to any natural object that is coloured. So although the masters do not have any visual experiences at all, it is claimed by Smart that they could come to hold for themselves the very same objective criteria for making colour judgements, and hence the very same concept of colour, that sighted people have.

Now one difficulty with this theory is that the masters in the parable would lack the direct awareness of colour which is customarily associated with the making of colour judgements. In some circumstances this would make their ability to judge less complete than that of normally sighted people. If, for example, there occurred a spectrum inversion, so that everything changed colour in a systematic way, but without upsetting the arrangement of the colour continuum, then this change would be undetectable to the masters, since they would be restricted in their 'perceptions' to the discriminatory responses of the slaves,

and these responses would not have changed at all. Ripe tomatoes and geraniums would still appear approximately the same colour as each other, but different from oranges and other things of similar hue; and oranges would still lie 'between' tomatoes and lemons in the colour continuum. This would be so even if tomatoes now appeared violet, oranges mauve, and lemons light blue to ordinarily sighted people. Smart does offer the suggestion that a general change of this sort would be detectable as a change in the wavelengths of light reflected by different objects before and after the spectrum inversion; but this is an ignoratio elenchi. For the spectrometer would indeed tell the masters that red things now registered the wavelength formerly registered by violet things, and so on. The masters could thus in a sense plot the spectrum inversion. However, there is a real sense in which the masters would still not know what change had occurred. They would know that objects had altered their properties in a way corresponding to a change in the properties whereby their slaves sorted piles of wool; but they would not know what we know when we say that an object has altered its standard appearance in respect of its colour. For the very notion of things having appearances would be quite foreign to the masters' understandings. Yet to say that one could have a concept of colour without having any idea of what it is for an object to present the appearance of being a colour is to reduce the concept of colour to something purely formal. Paraphrasing Kant, we might say that concepts without sensory content are in this case blind indeed. A similar point is made by Prof. D.W.Hamlyn in a discussion of the possibility of coming to know a particular colour via knowledge of its wavelength.* Hamlyn makes the point that such knowledge is inferential and the concept of the colour correspondingly incomplete until that knowledge is linked up with the ordinary visual perception of the colour; for there no inference is required. This, of course, is a step the blind masters in Smart's theory are unable to take.

This objection, though damaging in itself, points to a deeper difficulty with Smart's account. Those of us who

* 'Seeing Things As They Are', an inaugural lecture delivered at Birkbeck College, London (1965) pp.8-9.

are normally sighted can see colours, and can tell a blind person what colours things are. So we can fairly easily imagine a blind person learning about the colours of things in this way, and thus becoming able to converse fairly naturally about colours, either to other sighted persons or to other blind persons who had learned about colours in the same way. We could even imagine two blind persons agreeing that a certain room that they knew to be painted green was 'a restful green', judging this to be the case from the ways they had noticed sighted people behaving when in the room and when out of it. We might still feel, with some justification, that the expression 'a restful green' could not possibly mean to the blind persons quite what it means to us, who can experience the restful quality and what causes it. But we can envisage quite a substantial role for colour words to play in the discourse of the congenitally blind, always provided they gained their initial insight into colour language from the normally sighted. And this situation is just sufficiently close to the one Smart has described to mask the crucial difference. For in his account it is essential that the masters form the concept of colour for themselves, and not vicariously by taking over a range of concepts borrowed from the normally sighted slaves. Here it will not do to acquire, simply by different methods, a knowledge of the discriminations the slaves make in accordance with their concepts of colour. Nor will it do simply to instruct the slaves to sort the pieces of wool in 'some significant way', without giving them any indication of what way that is, and thus leaving it to the slaves to impose their own conceptual structure on the perceptions which determine this activity. For apart from the fact that the masters will simply have no conception of what the slaves are doing, even if the slaves themselves know, the strict requirement of Smart's position is that the slaves do not themselves know on what basis they are sorting the wool. This is the point of the taboo against the slaves using colour words when they speak to each other. The theory thus requires that the slaves just intuitively recognise for themselves that colours are colours, and a significant feature of things quite distinct from the other

significant features things may have in common. And the only way this could happen would be if the slaves were somehow immune to the force of the Private Language Argument, and could turn their own inner experiences into private exemplars for concepts which they could then consistently employ in their sorting activities. It is, of course, very far from Smart's intention to involve inner experiences in his theory in this way, or indeed in any other way. But the theory he presents absolutely requires it. It just will not do to introduce a discriminatory process without a conceptual structure to give that process sense; nor will it do to imagine a conceptual structure for the blind masters without perceptual content to give the concepts points of reference. For the former is incoherent and the latter empty.

Against this conclusion, it might be suggested that if we treat Smart's theory as a theory only about the ability to form colour concepts on the basis of discriminatory responses alone, then perhaps we can envisage the slaves (and through them, the masters) acquiring some idea of colour concepts without there being any opportunity for them to use colour words and acquire a notion of their right use in judgements. We might, for instance, imagine the spoken language being sidestepped in this instance, and replaced by some non-verbal system of signs for colours. However, this would still count as part of a language, because such signs could still be correctly or incorrectly used. There could still be rules for their application, then, and these rules would give sense to the otherwise unintelligible experiences that the slaves were having. However, I take it that Smart would want to exclude anything that led to the slaves having a conceptual grasp of what they were doing, so that really the masters' conceptual grasp of colours was parasitic on that of their slaves. So this way out is closed for his theory. And so, I believe, is the notion that the slaves might just gain a conceptual grasp of colour by analogy with other perceptually-orientated concepts which they are allowed to form by the usual process of coming to make sense of their experiences via an interpersonal conception, embodied in language, of how it is

normal to judge. It is by no means easy to see how this could occur by analogy without something central to the understanding of this sort of concept becoming lost or distorted. But in any case, Smart's theory is supposed to prove that the perceptual element can be replaced by a discriminatory-response element in every case within the theory of knowledge where there is a reference to these unscientific phenomena. So this effectively removes any role analogy might usefully play.

Smart has failed to furnish us with an adequate model of concept-formation divorced from sense experience of the appropriate kind. Does it follow that sense experience is necessary for anything which would count as a proper understanding of a concept which is sense-related, as opposed to the acquisition of a mere formal understanding of a concept derived vicariously from the appropriately-formed understanding of other people, as described in the previous paragraphs? I believe it does.

If one makes any intelligible judgement, one is exposed to the requirement that one should be able to justify one's judgement in some appropriate way. In the case of judgements involving observation, or derived from such judgements, and where it is claimed that the judgement is objective, we have already seen that the possibility of justification resides in agreement with other competent judges, and this in turn rests finally in the general agreement that we apply the concepts we use in judging, in a certain way which counts as the normal way to judge. But the justification for a subjective, or only putatively objective, judgement must be rather different. For there it need not be the case that one's judgement will agree with the judgements others make; and so the only justification which can in the end be offered (given that one claims to understand what one is saying) is that things simply do appear thus and so to one. And there, lack of the appropriate sensory faculty leaves no room for the notion of anything appearing thus and so to one. To point to this lack of a sensory faculty in someone is a decisive rebuttal to any claim he may make about how things appear to him in respect of that faculty. For to point to the

lack of a sensory faculty is to point to a reason for lacking the appropriate type of experience. Of course, the notion of something appearing thus and so also stands in contrast to the notion of something really being thus and so, which the idea of objectivity enshrines, as we saw in Chapter Three. This dual function of the word 'appears' stems from the fact that appearance, being the original basis for all sense-related judgement, is necessarily the original, and sometimes the only, basis for making a particular subjective judgement. However, while appearance may be contrasted with reality, it is equally true to say that the fact of things appearing thus and so is a necessary condition for any judgement, including the ones which describe how things really are, to be made. If things did not present appearances to one, then neither the judgement that they seem thus and so nor the judgement that they are really thus and so would be supportable.

Now it might be thought that I am begging a question here by referring to types of experience appropriate to the making of certain sorts of judgement. For could not a person come to make judgements of a certain sort via some unusual form of sense experience? In Chapter Three I discussed the possibility of people coming to be able to identify colours by touch alone, and suggested that providing this process produced accurate results, we might want to extend our idea of what colour was in order to take account of this new faculty of judgement. However, the ability to judge in this way must, at first, be unusual; there would be only a limited number of people who were able to have that experience in that way, and so could come to a position where significant agreement in judgements made in that way alone was possible. Thus the expression of judgements made in that way would at first have to employ the concepts relating to judgements made on account of the normal form that sort of experience takes. As long as that is the situation, the unusual form of experience will be only inferentially related to the normal form of judging.

One cannot rule out a priori the possibility of wholesale changes in human modes of experience coming about in some way. But if this were to happen, new conceptions

of how it is normal to judge would have to arise, and the concepts we have would change accordingly. For this reason, if the identification of colours by touch became commonplace, what we mean by 'colour' would also change. But until this happens, our conceptual grasp of colour, involving as it does the specification of the sense by which it is normal to judge colours, will continue to specify vision as the form of experience appropriate to such judgements. D.W.Hamlyn has made the point that while what is to count as objective may change, there are limits to the changes that are possible, just because we are the sort of creatures we are.* It might be said that the limits Hamlyn envisages are only contingent ones. The consequence of there being limits, however, is that they become limits on what could count for us as a form of life that we could understand by sharing in it. And this in turn sets necessary limits on what could count for us as an objective judgement or an objective way of judging.

I have now argued that having the appropriate sort of experience is a requirement for making an empirical judgement. The question remains as to how experience is deployed in judgement, and especially in subjective judgement.

Experience obviously constitutes a ground for judgement, in the ordinary sense that it is because I have the experience I do that I judge as I do and not in some other way. But the experience I am referring to here is, so to speak, objectified experience inasmuch as it is an experience such as another could have. This is because (as I argued earlier in this chapter) the experience is expressible, even to myself, only in terms that make sense; that is, only in terms derived from the concepts which I share with others. Wittgenstein regularly stresses the role of learning as a process of coming to understand via the public language and the forms of life that lie behind it; and he insists that in learning about a type of experience one does not simply come to attach a name to an experience one already recognises, but learns to identify and structure

* 'Objectivity' in 'Education and the Development of Reason' ed. Dearden, Hirst and Peters (Routledge 1972) p.256

one's experience according to the rules enshrined in public agreement. Such experience is then bound, in general, to be in accord with that of others. But the experience is still mine. That I have a particular experience does not entail that anyone else should have a similar one. (There may, after all, be nobody else about!) Now I may of course happen never to have recognised this type of uniqueness which my experiences have. Very young children presumably do not. But insofar as I come to be able to recognise the potential for this sort of uniqueness in my experience, the concept of subjectivity comes to have a point for me, and thus I come to be aware of myself as subject.

Wittgenstein makes just this point in 'Zettel'.* In ss423 he suggests that, once a child has learnt to use the word 'red', so that he can make all the usual judgements about things being or not being red, it is a further step for the child to learn of the red visual impression. This is because understanding the expression 'this seems/appears red to me' involves more than is involved in understanding the expression 'this is red'. The difference is that the latter expression lacks the sense of the observer as subject, while the former has acquired it. (Wittgenstein does not mention the intermediate category of objective appearances which I explored in Chapter Three. Consideration of this category would not alter his point, however, though it might complicate it; for objective appearances do not require the sense of the observer as subject.) Now because 'red' is a word in the public language, and redness an interpersonal concept, the notion of seeming red to me is parasitic on the notion of being red, and not vice versa. There is no risk, then, of our returning to an empiricist foundations-of-knowledge position in which being X is supposed to be entailed by seeming-X-to-me. The fact that I can now talk about how things seem or appear to me does not mean that how things seem or appear to me need be the source of justification for my judgements about the world. But it is a source of justification which is available to me (because I have a grasp of the concepts

* ed. von Wright and Anscombe (Blackwell 1967) ss418-35

and the language involved) in just those cases where disagreements arise. Introducing the concept of subjective appearance thus gives language a new joint, as Wittgenstein puts it there.

In view of all this, we must consider why Wittgenstein is sometimes held to have denied that sensations can be objects of knowledge for the person having them.. For if they cannot be objects of knowledge, then it would seem that they cannot be grounds for judgement for the person having them. And this would imply that I cannot treat the inner experience I have now as my reason for making a judgement which is ostensibly about that very experience. Now this contention seems wildly opposed to common sense. After all, it is surely on account of the pain I now feel that I utter the words 'I am in pain', just as it is on account of the pain that I might cry out. And in both cases it seems perfectly natural to say that I know why I made the utterance I did.

Nevertheless, this has been denied by some commentators on Wittgenstein's work, who maintain instead that at least some types of first-person present tense psychological sentences — the so-called avowals — can be neither true nor false, but should be assimilated to behavioural manifestations of the speaker's psychological state, on a level with cries, winces and the like.* But there have also been strong objections to this view. P.M.S.Hacker, for example, finds this an absurd position.** To mention only one pernicious dilemma that Hacker draws attention to, either memory reports are unknowable too, which means that it would be impossible for a person to remember how the pain he had yesterday felt, or else a sentence like 'this is the same pain I had yesterday' implies that I can know that a past pain is identical with a present pain, but without knowing that I have a pain now. Neither of these alternatives is very palatable; for surely I can remember a

* see N.Malcolm 'Wittgenstein's Philosophical Investigations' in 'Wittgenstein: the Philosophical Investigations' ed. G. Pitcher (Macmillan 1968) pp.77-83, and also P.T.Geach 'Mental Acts' (Routledge 1957) p.121. Geach supports the doctrine of avowals with some reservations.

** see his 'Insight and Illusion' (Oxford 1972) pp. 250-277

past experience and can compare it with a present one that I know I am having.

Another difficulty worth mentioning is that when we move away from the example of pains to less dramatic inner experiences, there seem to be precious few avowals for which there are any behavioural manifestations of the speaker's psychological state to which they may be assimilated. Perhaps it is thought that this does not matter, for we might regard the language of avowals as new (and more sophisticated) behaviour. But here the notion of behaviour is being unduly stretched. We should have to regard the avowals 'I'm seeing red' and 'I'm seeing green' as two distinctly different sorts of behaviour, whereas it seems wholly natural to regard both utterances as the same sort of behaviour, but relating to different experiences. And that would make the experiences something distinct from the behaviour, which the interpretation we are considering cannot allow.

Nevertheless, some remarks in the 'Philosophical Investigations' do appear to lend support to the view that Wittgenstein believed avowals to be non-cognitive and without truth-values. Malcolm (op.cit.) quotes with approval from ss244, where Wittgenstein wrote: "Here is one possibility: words are connected with the primitive, the natural, expressions of the sensation and used in their place. A child has hurt himself and he cries; and then adults talk to him and teach him exclamations and, later, sentences. They teach the child new pain-behaviour...the verbal expression of pain replaces crying and does not describe it." (Malcolm, incidentally, omits the first four words of this quotation.) Other remarks in this part of the 'Philosophical Investigations' might bear the same interpretation, notably ss246:
"It can't be said of me at all (except perhaps as a joke) that I know I am in pain. What is it supposed to mean — except perhaps that I am in pain?"*

I do not propose to attempt to give an exhaustive textual elucidation of Wittgenstein's doctrine of avowals, if indeed there is such a doctrine to be elucidated. I

* see also ss247 and ss288

only want to argue against the view that one's inner experiences cannot be objects of knowledge. For if this view is correct, a person cannot say that his experience is thus and so and have any reason for saying it if what he says is at variance with what everyone else says. Now we have already seen that an adequate account of objectivity must leave room for people to judge at variance with the majority. But there must be a reason for the individual's saying one thing rather than another when he does judge at variance with the majority view. And that he judges as he does because he has _this_ inner experience rather than _that_ one is the obvious explanation for something that _must_ occur if there is to be talk of subjective judgements occurring at all. We must try to see, then, what Wittgenstein is really denying when he denies that I can know that I am in pain; for it is upon a misunderstanding of such denials that the doctrine of avowals rests.

I think that only two points need to be made here. First, it must be borne in mind that when Wittgenstein denies that I can know my sensations, he does so in the context of an attack on private languages, and in opposition to the view that the _whole_ business of knowing what sensation I am having is a matter of private introspection. And this is opposition to a _theory_ which appears to treat private sensations as _the_ source of knowledge, but cannot do so because it only makes sense to speak of knowing where it also makes sense to speak of doubt, and of learning, and of all the other features appropriate to the public understanding of concepts and their right and wrong application to cases, which we have already examined. But the individual's relation to his own inner experiences is not of this kind. He does not know them _in this sense_ — that is, it makes no sense to talk, for instance, of his _coming_ to know them — because he _has_ them. Wittgenstein denies the appropriateness of the word 'know' here because he denies the appropriateness to the case in question of the language-game in which 'know' plays a part. The doctrine of avowals rests, I think, on the mistaken assumption that this is the _only_ way in which the word 'know' can be used.

Wittgenstein has been interpreted as making a categorical denial where, as so often, he was only issuing a warning against using a certain type of expression because it tempts us into a particular way of thinking when doing philosophy. He was not denying that the relation between an individual and his inner experiences could involve awareness and indeed recognition (once the concepts and the public language that expresses them have been jointly mastered). He was only denying that our understanding of our experience begins with a process of private introspection that looks as if it could amount to knowledge on its own.

Because the word 'know' is being discussed in the context of its use in a private language, then, it is that use which is being outlawed. But there is a quite different sense in which we can speak of a person knowing he is having a particular sensation. For he can be justified in what he is saying when he reports a particular sensation, or when he just thinks the words of that report to himself. The justification here is that he has mastered the language — that he speaks English. Now the proof that he has mastered the language is not, as is often supposed, that he always says, or is inclined to say, what others would regard as the appropriate response in the circumstances, but only that he normally does. We are quite used to the idea that someone can be sincere in saying that he feels no pain when his behaviour or circumstances suggest otherwise; or that a person can say and mean that something looks orange to him when others would say it was purple. Here we look for a physical explanation if we look for anything. Only if the deviant judgements persist in the face of a lack of any recognised physical explanation do we doubt the person's grasp of the language or the concept involved in this sort of judgement. When a person does understand the language, then, what are we to make of the fact that on occasion he says something different from what others say in the same circumstances? Instead of assuming, as the behaviourist must, that he has temporarily lost his grasp on the concept involved, we are quite entitled to say that he has retained that mastery and knows what he is saying. And if he knows what he is saying, then he knows what experience he is

having. There is nothing wrong, then, in saying that a man who reports a subjective experience knows why he says what he says. He says what he says because he knows what the experience he is having is like. And this is all that is needed to make the notion of a subjective judgement being based on knowledge of one's own experiences an intelligible and coherent one.

In this chapter I have tried to draw out from the work of the later Wittgenstein those elements which, taken together, provide us with an explanation of how our most basic empirical judgements can be objective whilst at the same time the possibility of meaningful subjective judgement is preserved. For this account explains why people must normally agree in their judgements about the world, yet also allows that they are judging on the basis of their own experiences. By making normal agreement in judgements a precondition of the intelligibility of one's own experience, but yet not a substitute for it, we are enabled to see why people must as a rule agree in what they say about the world. The individual's judgements, then, will normally be made with justification; and when two further conditions are met, that justification will be sufficient for the judgements to qualify as objective. The first of these conditions is that the judgement should accord with those made by other people, which of course it normally will. But to avoid allowing this analysis of objectivity to lapse back into a consensus theory, we must also add a second condition, which is that on this occasion other people are judging with an understanding of what they are about. For if they are not judging with understanding, then the possibility will arise (as sometimes happens) that the individual <u>is</u> being objective in his judgements even though his judgement is different from the judgement of the majority in this case. Such cases are necessarily abnormal, because if that were not so, we should have to say that the majority was judging without understanding. While that might happen in a particular area of judgement, we could not expect it to happen generally without also having to say that here human understanding had broken down altogether, if indeed it had ever come to exist in the first place. It is because

these cases are abnormal that we look for explanations when they occur, and do not just accept them as matters of personal opinion in an area where judgement is inappropriate. For example, if we imagine that a group of travellers in a desert all see an oasis on the horizon, they might all estimate that it can be reached in an hour or so — all except their guide, who understands how to judge these things, and knows a mirage when he sees it, and can estimate that it is three days' journey to the oasis. Here the majority is wrong because the majority lack a proper understanding of how to judge such things. But if instead we imagine that the desert is a place where nobody, however practised, can judge distances accurately, where no two people ever agree on how near or how far anything is, and where nobody ever judges right (for in this case we have an independent way of finding out the truth, by measuring the time taken to reach a point at a steady speed); then here we must say that there is no understanding of how to judge distance by eye alone, something people can do as a rule — but not in the desert. A more homely example of this lack of ability to judge being based on a lack of understanding occurs in the well-known experiment in which a group of young children all say that when a quantity of water is poured from a squat vessel into a tall thin vessel, there is more water in the second vessel than there was in the first one. Here it is obvious to us that the children have failed to judge objectively, not because their normal understanding of how to judge quantity has been distorted by factors that it is impossible to take account of, but because as yet they lack an adequate understanding of the concepts involved in the judgements they are called upon to make.* We must bear in mind that the majority may always risk failing to be objective for either of the two reasons suggested by these examples, a failure to understand the relevant concept involved, or a failure of experience or imagination which would allow them to apply the under-

* see J.Piaget & A.Szeminska 'The Child's Conception of Number' (Routledge 1952). For an extensive bibliography of this topic, see D.Child 'Psychology and the Teacher' (Holt 1973) pp.90-93.

standing they have to the case before them. Both factors can apply to individuals, to groups, or to whole societies. While this does not often happen in the cases of simple judgement we have mainly been discussing, it does increasingly as the matter to be judged involves more complex chains of inference or more detailed marshalling of the facts involved, or where the judgement is made in accordance with a body of theoretical presuppositions, as happens in the sciences particularly. But to pursue this line of enquiry would take us too far from the theme of this work, which is to trace a link between the conditions for the objectivity of everyday ascriptions of quality in empirical judgements and those for basic ascriptions of value.

The argument in this chapter has centred on the individual's ability to judge objectively in matters in which his sense-experience plays a crucial role. However, the conclusions we have reached, and the argument which has been vital to our ability to reach them, has given us the key to much more than this. For it has shown us the general connection that exists between understanding and agreement in judgements. If an idea can be expressed and communicated, it must be deployed according to the rules that make it intelligible in the public language. And wherever the possibility of a justification for a judgement exists, such justification must take the form of there being, at some level, agreement in the results of our judgements. We have seen how even the justification for having an experience which others do not have turns on the claim to have a mastery of the concepts involved in the judgement; and this claim is in turn justified by the fact that normally the person making the claim will judge as others do. When Wittgenstein established the point in the'Philosophical Investigations' ss242 that understanding involves there being agreement in judgements, there is no reason to think that he was only talking about judgements which are sense-related, or even that his remarks were confined to empirical judgements. His example of measuring is hardly a case that would spring to mind as typical of judgements involving inner experience. And though we would perhaps think of measurement as first and foremost a variety of empirical judgement, it is not

far-fetched to regard measurement as a process with many points of similarity to evaluation.* Evaluation is often portrayed as if it were a species of liking or preferring. Even here it is possible to misuse those terms in such a way as to show a lack of understanding in what is involved (as when a man claims to prefer everything, or to like what he always shuns without cause). But as soon as any substantial mechanisms of interpersonal justification, proof, reasons or evidence come into play, we can see clearly that the claim to understand what one is saying when one utters a judgement is firmly linked to the idea that one judges normally in such cases. The more substantial these mechanisms, the firmer is the link with normal judgement. To judge normally involves judging for the most part as others do; and it implies that objectivity has a place in that type of judgement. These are issues to be examined in more detail in later chapters. But in turning more directly to questions of evaluation now, we will begin by challenging an assumption about the way in which facts and values would supposedly have to be related if evaluative objectivity were to be possible. This is almost the last line of defence that the sceptic about moral objectivity has; and we need to probe these defences before we can do justice to the questions of moral evidence and moral meaning which I have begun to raise here.

* c.f. J.O.Urmson 'On Grading' (Mind 1950)

CHAPTER SIX

THE FACT-VALUE DISTINCTION

In Chapter One of this thesis, I spoke of there being three major and to some extent related misconceptions about the nature of objectivity which, between them, so obscured the question of what objective evaluations would be like as to make it impossible to discuss whether any evaluative judgements are objective without prejudging the issue. Two of these misconceptions have now been scrutinised and from the discussion a more satisfactory analysis of what it means to call a judgement objective has, I believe, emerged. It remains in this chapter to consider the third area of misunderstanding, the notion that there is a fundamental difference between facts and values, and that that difference has the consequence of rendering evaluative objectivity impossible. Sorting out this muddle will lead us directly to the first part of the account of the possibility of moral objectivity and of the extent to which it exists.

All three misconceptions about objectivity to which I am referring have this in common: they stem from an outlook which regards objectivity as essentially something to do with objects, and usually material objects at that. Only those characteristics manifested by judgements about objects, it is assumed, can have any claim to objectivity; and this comes to mean that being factual and being objective are treated as identical states. The first of these misconceptions, it will be remembered, was that to call a judgement objective was in effect to say that it was true; and, moreover, to say that it was true in a way that, on a generous interpretation, empirical judgements are true — that is, true of a reality existing independent of the individual's mind. This assumption was, however, seen to be false, because appearances can themselves be objective, and because the nature of empirical knowledge is such that there is a dynamic interaction between objectivity and truth which permits each to modify the other in a way that would be impossible if the two were identical. This disposed of the assumption that judgements could only count

as objective if they were true just as empirical judgements are. The second misconception involved a fundamental mistake over the reasons why a judgement about one's experience could be called objective. It was supposed that because the origins of one's experiental judgements lie in one's own capacity for having sensations, the logical basis of claims that one's judgements about the world are objective must somehow lie there too. But this is entirely wrong. As I have argued, our capacity to call a judgement of this sort objective indicates that there is a public standard with which our individual judgements must normally accord, and which in fact renders our individual experience intelligible. This is necessarily so, because the language in which such judgements are expressed can be meaningful only if there is agreement on how such judgements are to be properly made; and this, we saw, imposed a structure on judgements in general. Although the judgements so far discussed have been judgements made on the basis of sense-experience, this was only because they are typical of objective judgements, and exemplify their properties clearly. It is not, as the following arguments will show, because objectivity is in any way logically restricted to the sorts of judgements which can express facts.

The third and last of the misconceptions which make it so difficult to discuss moral objectivity at all adequately is generally known as the fact-value distinction. It can be presented in the form of several different arguments, and three well-known forms of it will be examined in due course.* As well as appearing in a number of guises, the argument has been used to prove different things — that arguing from 'is' to 'ought' is illegitimate, that morals are autonomous, that moral language is radically different from descriptive or factual language, and of course that evaluations cannot be objective. Some of these issues, and the question of whether the fact-value distinction proves what it purports to prove, are not my immediate concern. Adequate discussion of them would take us too far from the present line of argument. I am inclined to think that the

* by Hume, Toulmin and Moore.

argument is generally suspect, because in all its forms it turns on playing up one apparent difference between empirical and evaluative judgements, and insinuating that this is a crucial difference — which it is not. In criticising the fact-value distinction, I am not of course implying that there are no differences at all between the language of facts and the language of values, or that for example judging that something is yellow is just like judging that something is good. I am only criticising the interpretation of one particular point of superficial dissimilarity, and the undue weight that it has been made to bear. The belief that this dissimilarity is crucial stems as before from the mistaken view that ordinary empirical judgements are in themselves objective, and thus represent the only way in which objectivity can be manifested.

For the purposes of discussion, I shall first present the argument for the fact-value distinction in the simplified form in which it appeared in Chapter One, and which directly attacks the notion of evaluative objectivity. It runs as follows:

Ordinary empirical judgements, such as those based on direct observation, are always right or wrong. We may not always be able to decide, in a particular case, whether a judgement is correct or not, but even then we can say in what circumstances we could determine this. So there is no room for significant and indeterminable disagreement in the realm of ordinary empirical judgement. Since such judgements can be shown to be either right or wrong, empirical judgements which are made correctly count as objective. However, evaluative judgements are completely different in this respect. People do disagree widely in their evaluations, even when there is no disagreement between them on matters of fact. And it is always logically possible for two people to agree on every point of fact and yet arrive at different evaluative conclusions. Thus the notions of correct and incorrect judgement cannot apply to evaluative judgement, and consequently such judgements cannot be objective.

As I have presented it here, this argument is defective in three distinct ways. First, it contains a doubtful

factual claim; second, it unfairly compares an epistemic point with a logical consideration; thirdly, it makes an important epistemological assumption. Not all these errors occur in every variant of the argument, and other mistakes are sometimes present instead. I shall draw attention to a fourth one later, but for the moment we have enough to be going on with.

The dubious factual claim in the argument is perhaps the hardest to spot, just because the very pervasiveness of this type of argument has conditioned us against recognising it. This is the belief that there is much more disagreement amongst people's evaluative judgements than there is amongst people's judgements about the contents of the material world, other things being equal. But is this really so? To take a simple example: if this claim were true, we should expect there to be much more disagreement about whether sharpness was a desirable quality in a knife than about whether grass is green. But in fact there is universal agreement on both these points. When we think about disagreements in evaluation, it is almost always the case that we think of contentious moral or aesthetic questions which are, as often as not, very complex. When we think of factual judgements, on the other hand, we tend to think of observational judgements, or judgements which at best involve one or two simple inferential steps. To make a fair comparison with a problem of moral complexity, we should perhaps imagine a medical student who has been confronted by a consultant with a patient displaying a range of symptoms not usually found together, and being asked to state what disease the patient has in fact got. The assumption that an evaluative judgement, of whatever sort and in whatever circumstances it is made, is bound to present diverse results, whereas empirical judgement will present uniform and easy conclusions, begs far too many questions for comfort or accuracy. Suffice to say here that it predisposes us to accept that we are facing two radically distinct types of judgement.

We may perhaps shrug aside this point as no more than a piece of simplification which is unavoidable in a general argument to which it is not in any case essential, since

the argument really aims to make a logical rather than a factual distinction between the two types of judgement it deals with. (We might wonder, though, whether the argument would be quite as attractive if, say, moral disagreement were as rare as disagreement about the colour of grass.) My next objection, however, is more radical. The argument we are considering begins by telling us that there exist decision procedures for deciding on the correctness or otherwise of empirical judgements in all cases. (This is not quite true, in fact, even if we add an 'in principle' clause. There is always the possibility that we may come up against novel species of experience which no existing decision procedures can cope with. At least I do not see how we can be sure a priori that we will never have to find new ways of deciding what is the case. The 'in principle' clause will have to be streched very taut to cover these cases as well. This is not a central point; but it does no harm to remind ourselves from time to time that the contrast being drawn is not as simple as it looks.) If the existence of decision procedures is being emphasised as a feature of empirical judgement, then, we might expect that the point of contrast with evaluative judgements would consist of giving a reason for saying that decision procedures are lacking for some, or perhaps all, evaluative judgements. But the argument does not do this. Instead, it purports to show that there could not be decision procedures for evaluative judgements. (This is a simpler, if more radical, step than attempting to argue that some central class of evaluative judgements involves a feature that renders it indeterminate.*) What is supposed to show the impossibility of there being decision procedures for evaluative judgements, and hence of there being evaluative objectivity, is the fact that it is logically possible in any case of evaluation for people to disagree. I do not have any hesitation in calling this claimed logical possibility a fact. For of course this logical possibility must exist. If it did not, the question of the objectivity of evaluative judgements could not be asked and would not be argued over. The very same consideration applies also to empirical

* as Bernard Williams has tried to do. See Chapter Seven.

judgements, of course. Without the logical possibility of there being disagreements in empirical judgements, there would be no decision procedures for settling such disagreements, because there would be no need for them; and one half of the original contrast offered by the fact-value argument would have disappeared. Furthermore, since decision procedures are available for, and applicable to, all or virtually all empirical judgements, it must follow that the logical possibility of disagreement is similarly present — as of course we know it is since it is a cardinal feature of empirical judgement that it is not a tautology.

So far, then, this supposedly damning feature of evaluative judgement appears to be equally, and necessarily, present in the empirical judgements which were being contrasted with evaluative judgements in this very respect. Any evaluative judgement can, logically, be the subject of disagreement, but then so can any empirical judgement. Perhaps, though, we are really being asked to envisage a rather different possibility. The argument might be asking us to consider the possibility of of there being disagreement about all evaluative judgements. And this is something one can, I suppose, imagine. Perhaps the argument is asking us to contrast this possibility with the state of affairs which might exist with regard to empirical judgements. For since there are decision procedures for settling empirical disputes, how could there be disagreement in every case of empirical judgement? We must remember, however, that the contrast presented by the fact-value argument is a logical one, or purports to be so. The proper contrast is then between the logical possibility of disagreement in every evaluation and the same logical possibility in every judgement about what is the case in the world. We have seen already that this logical possibility exists for any empirical judgement; it must therefore exist for all. To be sure, if it were actually the case that nobody agreed with anybody else on any question of fact, that would show that there were not really any decision procedures after all; and the argument we are discussing begins, quite correctly, by asserting that there are. But it is no more difficult to imagine, as a logical possibility, that nobody agrees on any factual question with anybody else, and that

what seemed to be ways of settling factual disputes no longer command any respect or carry any conviction, than it is to imagine that nobody ever agrees on a question of values. In the way in which the idea of decision procedures is referred to in the argument, however, there is a confusion between the fact that there are empirical decision procedures and the suggestion that this makes the possibility of universal disagreement over matters of fact a logical impossibility. Once again, what is the case has to be conflated with what might be the case to make this reading of the argument a credible one.

So on either reading of the argument in question, the impossibility of there being ways of settling evaluative disputes has not been demonstrated by drawing attention to the possibility of disagreement. What makes the argument superficially plausible is that, as I have suggested, we are inclined to assume that there is a difference in practice between the inherent decidability of empirical claims and the inherent undecidability of evaluative ones. And if we are already inclined to believe this, it is comparatively easy to give the impression that there is a logical difference here, by making a logical point about one side of the dichotomy but not the other, and thus giving the impression that a contrast of significance has been drawn.

I turn now to the third difficulty with the argument, which I referred to earlier as an epistemological assumption. As I have presented it, the fact-value argument emphasises the decidability of factual questions and contrasts this with both the (dubious) fact of widespread evaluative disagreement and the (wholly unproven and untenable) theory that this disagreement is somehow a logical characteristic of evaluative judgements but not of empirical ones — so implying that such disagreement indicates undecidability. No mention is made in the argument that there might be decision procedures which apply specifically to evaluative judgements. It is merely assumed that none such exist. And this assumption is made plausible by giving the impression that a type of decision procedure has been tested on evaluative judgements and found to be unsuitable. The further assumption that no other decision procedures could

exist for deciding between rival evaluative judgements is then subtly implied. The decision procedures offered for evaluative judgements are, unsurprisingly, exactly the same ones as apply to empirical judgements. They are whatever methods enable us to determine what is the case, to decide between rival empirical claims. For it is implied that any objectivity that any evaluation might possess could come only from the facts on which the evaluation is based; and that the possibility of there being valid decision procedures ceases at that point. We can justifiably claim to know what the facts <u>are</u>, but not what the facts <u>show</u>, it is suggested. For, we are told, it is possible for people to disagree over the evaluative conclusions they reach even though they are in complete agreement over 'the facts'.

I am not of course denying that evaluative conclusions are reached by drawing inferences from 'the facts'. This is sometimes expressed by saying that evaluative terms are supervenient. But here we face the question of <u>which</u> facts are relevant, and <u>what</u> conclusions are to be drawn from them. And it is just assumed that there is no way of determining this. The stranglehold of the procedures for deciding what is the case in empirical matters has been so firm upon the philosophical imagination that only two possibilities have seemed viable. Either moral judgements are determinable in exactly the same way as empirical judgments are, because evaluation is a process exactly like empirical judgement — this leads to Ethical Intuitionism — or else what is the case in the world must of itself entail unequivocal moral conclusions, so that the objectivity of moral judgements would be a logical consequence of the objectivity of factual judgements, facts being the place where objectivity has its proper dwelling. If neither of these possibilities proves acceptable, it follows that evaluations cannot be objective.

But it is absurd to suppose that these are the only possibilities. If people disagree over what is good or what should be done although they agree over 'the facts', it is evident that their disagreement is not about what the facts are, but about which of them are relevant and perhaps decisive for the evaluative question that people face. And to say that there is disagreement here may be to say no

more than that people are liable to misjudge, especially if the case is complicated. If people agree, not only on 'the facts' but on how the facts constitute evidence, and for what, it would be absurd then to suppose that disagreement would be endemic without the presence also of widespread error. The fact-value argument invites us, however, simply to ignore such considerations. It has laid the ground-rules for the debate. Moral conclusions are to be deduced from consideration of the facts alone, or they are to be denied any status higher than subjective opinion, and are to be regarded as strictly groundless.

The outlook underlying the argument I have been discussing, with its emphasis on the special role of facts as both the exemplars and the source of objectivity, has had a profound influence on English moral philosophy. Both those who have sought an objective basis for moral principles and those who have denied that this is possible have been conditioned by the ideas that there is a major difference between factual and evaluative discourse, that this difference is essentially a logical difference, and that the logical difference has consequences for the possibility of justified beliefs about value. I have indicated that there is reason to doubt all this. Yet the influence of this outlook has been pervasive. It distorts our view of a crucial passage in Hume's 'Treatise'; and G.E.Moore, who maintained that moral discourse could be objective, nevertheless constructed large parts of his moral theory under the influence of this outlook, on a direct analogy with empirical observation statements. The influence of the outlook was a distinct weakness in the whole approach of Ethical Intuitionism, and consequently on its later opponents such as Ayer, Stevenson and Hare. Perhaps, however, the clearest example in the recent literature of the continued impact of this outlook on moral philosophy is to be found in the opening chapters of Stephen Toulmin's 'The Place of Reason in Ethics'.* In the first part of his book, Toulmin examines and rejects three approaches to ethical

* (C.U.P. 1950)

theory which he calls collectively 'the Traditional Method' and the aim of which he describes as "not so much to discover what reasons and arguments should be accepted in support of ethical decisions, as to pin down — to characterise — ethical concepts by means of some kind of definition" (p.5). The first of these approaches to ethical theory is labelled 'the Objective Approach', which Toulmin describes as "one of the oldest and most familiar doctrines of philosophical ethics" (p.9). This is "the doctrine that, in saying that anything is good or right, we are mentioning a property which it has, the property of goodness or rightness" (ibid.). And it is quite clear, although Toulmin does not say so, that he takes any theory which claims objectivity for moral judgements to amount to this, since he thinks that only judgements about properties are capable of being objective. Toulmin then goes on to construe properties as being characteristically the features possessed by material objects (his two main examples of a property being redness and 259-sidedness) and constructs a list of the sources of possible disagreement about whether something possesses such a property. The list has four main headings: deception, organic defect, linguistic and definitional differences, and incorrect application of an appropriate routine for deciding (such as measuring or comparing). Toulmin claims that his list is exhaustive, which in broad terms it is, and his strategy is then to argue that these reasons cannot account for the fact that there are moral disagreements. I shall argue later in this chapter that this is incorrect, because Toulmin simply ignores the possibility that routines for deciding could apply to any judgement which is not empirical. However, he is putting forward a standard argument against Ethical Intuitionism here, so he takes his argument so far as a proof that moral terms are not the names of properties.

That is fair enough. The concept of a property is so firmly annexed to the concept of a material object that it is unwise to use the term 'property' in connection with anything else. What is revealing — and objectionable — in Toulmin's argument is that once we have disposed of the idea that moral words refer to properties in just the

way that, say, 'red' does, we hear no more of moral objectivity. For the doctrine that moral judgements are objective is taken by him to be synonymous with the idea that there are moral properties on an exact analogy with the properties of material objects. Indeed, in diagnosing the mistake which advocates of property-objectivism in ethics have made, Toulmin correctly suggests that the model of the objective property belonging to a material object, which is the subject of discussion in a dispute about whether something is, say, red, "is so compelling" that the urge to treat moral disputes as being just like that "is almost irresistable".* However, he fails to appreciate that in treating objectivity as a characteristic exclusive to judgements about material properties and of any other judgements which correspond exactly to his account of how we decide that a material object has a particular property, he is himself being compelled irresistably by another aspect of the very same model into the assumption that there is no more to be said about the possibility of moral objectivity.

To be fair to Toulmin, he is doing no more than restating the usual sort of rebuttal to Ethical Intuitionism. He is not particularly interested in moral objectivity as such. But it is indicative of the prevalence of the misconception I have been discussing that he calls the doctrine of moral properties construed on the model of material properties the objective approach. This shows how far the ideas of objectivity and of judgements about material objects have become intertwined, to the detriment of the possibility of a proper examination of the question of moral objectivity. Yet, perversely enough, the origins of the fact-value distinction, in the celebrated passage in Hume's 'Treatise', concern a rather different distinction which Hume noted between factual and moral judgements; and it is Hume's interpreters who have forced the argument into its present form. The substance of Hume's argument in the first section of Book III of the 'Treatise', from which the passage comes, is that moral distinctions cannot be derived by reason alone, since reason of itself cannot move us to act; nor can they

* op.cit. p.28

be matters of fact discoverable by the understanding. If any fact is involved, Hume thinks it is the fact that I feel a sentiment of approbation or disapprobation when confronting a given state of affairs, a sentiment which here he rashly or carelessly seems to identify with moral judgement itself. His real point, then, in comparing factual and moral judgements, is that they have different origins. Facts, as Hume conceives them, originate from impressions, which at this point he is content to suppose indicate real existences external to the mind. On the other hand, "vice and virtue...are not qualities in objects, but perceptions in the mind".* Hume is therefore implying a distinction between facts and values rather like the difference between primary and secondary qualities of objects, but with the additional difference that these "perceptions in the mind" involve the will and engage the passions. This way of presenting the matter (which, characteristically for Hume, is a mixture of epistemology and classificatory psychology) simply brings out the fact that evaluative judgements involve different mental faculties from judgements about facts (or, as I would prefer to put it, that evaluation brings our aims or interests into the reckoning in a way that factual judgement does not). In pointing out, therefore, that when "every system of morality, which I have hitherto met with" begins with factual considerations and then slides into evaluation, Hume can be understood as pointing to the need to recognise that a new sort of justification is required for the evaluations. It will not do to suppose either that what is now being discussed is still a matter of fact, or that the facts on their own lead unequivocally to an evaluative conclusion unless further mental powers are brought into play — and therefore, I maintain, without a further level of justification becoming appropriate. So this section of the 'Treatise' is not an argument against the <u>possibility</u> of moral objectivity. It only becomes that if <u>we</u> add the assumption that only facts can be objective, and therefore that the only way in which moral objectivity could arise is by a process of entailment from 'the facts'

* 'Treatise' III,i,1. p.245 of Hume's Works (vol.2) ed. Green & Grose (Scientia Verlag Aalen 1964)

alone. Hume does not make this assumption because it would be superfluous to his overall argument. For since he has already shown that moral judgements involve the operation of the passions, whereas factual judgements do not, there could be no question for him but that these are two distinct ways of judging. Thus the question of whether the correct standard of judgement for the one type of judgement is also the correct standard for the other type simply does not arise. Hume is led towards moral subjectivism* by a quite different consideration, his identification of moral judgement with the individual's feeling of approbation or disapprobation. This leads him into difficulties elsewhere in explaining why people commonly share the same sentiments, but sometimes do not. However, that is a problem we need not pursue here.

My claim, then, is that the original version of the fact-value argument, read in its context, does not aim to do more than point out the error of arguing from one kind of judgement to another as if there were no difference between them. That Hume finds "deduction" objectionable when anyone tries to do this is no more than we might expect from his sceptical approach to philosophy. He does, after all, indicate similar difficulties with the process of arguing from particular premisses to general conclusions. Are we to infer from that (pace Hume) that there are never good grounds for belief in a conclusion in a causal argument, or for preferring one causal explanation to another? (The possibility of giving a parallel account of moral judgements and inductive judgements is explored by A.C.MacIntyre in 'Hume on Is and Ought': Phil. Rev. vol.LXVIII,1959.)

What scepticism really does is to disbar certain sorts of justification and explanation from fields where they are held to be inappropriate; and it does this with a greater or lesser degree of severity. Only the most extreme sceptic would, however, claim that by doing this he destroys all possibility of explanation or justification except by de-

* or, according to A.G.N.Flew, towards emotivism. (see 'On the Interpretation of Hume' (Philosophy, vol.XXXVIII, 1963)). There is considerable debate about what Hume's moral philosophy amounts to. R.S.Peters regards Hume's notion of moral approbation to be sufficiently sophisticated to make him an objectivist without knowing it. (see 'Ethics and Education' (Unwin 1970) pp.108-111.)

duction from what can be known with certainty. Yet that, in effect, is what Hume's argument has been made to do in respect of evaluative judgement. 'Facts are objective; facts alone do not logically entail evaluations; hence there can be no adequate basis for any evaluative claim.' — that, in essence, is what many philosophers have made of Hume's distinction in relation to the question of moral objectivity.

It is worth comparing this with what Hume actually wrote in that famous passage. His words are: "I cannot forbear adding to these reasonings an observation which may, perhaps, be found of some importance. In every system of morality, which I have hitherto met with, I have always remarked that the author proceeds for some time in the ordinary way of reasoning, and establishes the being of a God, or makes observations concerning human affairs; when of a sudden I am surprised to find that, instead of the usual copulations of propositions is and is not, I meet with no proposition that is not connected with an ought or an ought not. This change is imperceptible; but is, however, of the last consequence. For as this ought or ought not expresses some new relation or affirmation, it is necessary that it should be observed and explained, and at the same time that a reason should be given for what seems altogether inconceivable, how this new relation can be a deduction from others, which are entirely different from it. But as authors do not commonly use this precaution, I shall presume to recommend it to the readers; and am persuaded, that this small attention would subvert all the vulgar systems of morality, and let us see that the distinction of vice and virtue is not founded merely on the relations of objects, nor is perceived by reason." (op.cit. pp.245-6).

Advocates of several moral theories have referred to this passage approvingly as lending support to their views. H.A.Prichard quotes Hume in support of Intuitionism, A.J. Ayer refers to Hume's discovery of a point of logic which supports Emotivism, and R.M.Hare refers to the logical impossibility of deriving 'ought' from 'is' as Hume's Law.*

* in, respectively, 'Moral Obligation' (Oxford 1949) p.89, 'Logical Positivism' (Glencoe 1959) p.22, 'Freedom and Reason' (Oxford 1963) p.108.

Hare uses Hume's argument in support of Prescriptivism. That all these philosophers take Hume's argument as supporting their own is understandable, since all their theories have the denial of Naturalism in common. Naturalism is the idea that certain states of affairs in the world indicate of their own accord how we ought to act; and Hume is undoubtedly opposed to one form of Naturalism, the belief "that the distinction of vice and virtue is...founded merely on the relations of objects..." (loc.cit.).

These philosophers have, I think, broadly agreed in taking Hume to be saying three things. First, that the type of moral argument he wishes to call into question has facts as its premisses and concludes with an evaluation; second, that the inference involved here, which Hume calls deduction, is a strict logical entailment; third, that valid arguments are either deductive or defective. From these three points they take Hume to be arguing that moral conclusions cannot be arrived at deductively, and hence cannot be susceptible to impersonal proof. Since 'the facts' cannot then furnish premisses for deducing a moral conclusion by themselves, and since it is assumed that objectivity could derive only from the facts and what follows from them, it is assumed that moral judgements could not be objective. Now <u>if</u> this is what Hume is really saying, then I have argued already that such a line of reasoning would entirely fail to disprove the possibility of moral objectivity, since it ignores the possibility of there being ways of deciding between contrasting moral conclusions that do not turn simply on our being able to make deductions from 'the facts', but which nevertheless guarantee correctness — because the moral conclusions involved may be drawn from the <u>relevant</u> facts, and because that process is one that can be <u>justified</u>. The false assumption is that there could not be such procedures for deciding moral questions, because objectivity can only be a feature of factual judgements and what is exactly like a factual judgement or can be entailed from a matter of fact. So even if we accept this interpretation of Hume's argument, it proves much less than it is thought to do with respect to the objectivity of moral judgements.

However, it is by no means certain that Hume is saying what these advocates of his views say he is. First of all, Hume claims to be putting foward an argument which "would subvert all the vulgar systems of morality";* and he also makes it clear at the beginning of the paragraph that he is talking about systematic justifications of morality. It was certainly customary in the Eighteenth Century and earlier for writers to attempt to derive particular moral codes from considerations of fact, whether natural or supernatural, or from speculations about the workings of the mind. But to say that moral judgements are not susceptible to proof if they are derived in this way is not quite the same as denying that moral judgements are without justification when they are made properly; nor indeed is it quite the same as denying that systematic proof might be forthcoming in some different way. (Prof. MacIntyre, for example, interprets Hume as believing that moral rules are justified providing that they are in everyone's long-term interest, and denies that this is, for Hume, a breach of his own argument.**) However, proponents of 'Hume's Law' have certainly assumed both that Hume was describing the way in which we actually make moral judgements, rather than the way in which a systematic moraliser might claim that we are justified in coming to make them; and also that any other way of justifying them has been by implication ruled out. And neither of these assumptions fits Hume's stated intention here. Hume himself may not have been clear about difference between a claim that a particular form of judgement is justified and the claim that it is related to other, supposedly more secure, types of justified judgement, tending as a sceptic to collapse the first claim into the second. His more recent interpreters have rather less excuse for failing to see the difference.

A second reason for casting doubt on the interpretation of Hume's passage which I am discussing, is his use of the term 'deduction', combined with the idea that arguments which are not deductive in the strict sense given to the

* my italics

** A.C.MacIntyre 'Hume on Is and Ought' (Philosophical Review 1959)

term in modern logic are defective and do not provide proof of the conclusion as they are required to do. When Hume spoke of deduction, he may have meant merely inference, i.e. any form of non-intuitive reasoning. Hume did sometimes use the term 'deduction' where we should nowadays speak of induction. If this is the case, then modern attempts to describe a process of moral argument as the entailment of evaluative conclusions from factual premisses, and then to deny that this can be justified, seem rather more remote from what Hume was saying. Prof. R.F.Atkinson has pointed out that Reid, Hume's contemporary, took Hume to mean entailment; but what Reid thought is less than conclusive.* If Hume found difficulty with the logical relation between factual statements and evaluative conclusions in a rather wider sense than is suggested by supposing that for him 'deduction' meant entailment, then this points to a general difficulty inherent in Humean scepticism, and not to a particular problem about one sort of relation. When Hume remarks "a reason should be given for what seems altogether inconceivable, how this new relation can be a deduction from others, which are entirely different from it", we must bear in mind Hume's belief that he has already shown that moral judgement involves the introduction of entirely new elements, the passions, into the argumentation. This is something which no inferential process, however broadly conceived, could account for by itself. Hence his demand for a reason, and some later attempts to supply it. This also explains Hume's use of the word "merely" in the final sentence of the passage I quoted earlier. For it is not how things are in the world alone which can lead to an evaluative judgement of any sort.

For these reasons, I take Hume to be saying something rather different from, and more general than, the view which the other philosophers I have mentioned have ascribed to him. He is pointing to a defect in one common way of philosophising, that of presenting values as systematically derivable from matters of fact alone, and calling for a general explanation of how such derivations can be reliably

* R.F.Atkinson 'Hume on 'Is' and 'Ought': a Reply to A.C. MacIntyre' (Philosophical Review 1961) p.235

made. As a sceptic, and given his belief in the involvement of the passions in all moral judgements, he is naturally disinclined to think that an explanation can be given in such a way. His interpreters, however, have seen in this argument a model of how ordinary moral argument is supposed to function, and have found that model mistaken — as of course it is. For no fact, or collection of facts, can by itself entail that something has value, or that I ought to do something. The danger occurs when this belief is combined with the further idea that only by such a process as Hume has described (or has been thought to describe) can moral and other evaluative judgements be capable of possessing objective justification. There is nothing in Hume's argument to suggest this further idea; but the whole temper of his sceptical approach has been taken as implying it. This being so, it is not surprising that a consequence often drawn from Hume's argument is that evaluative judgements are demonstrably non-objective; and the demonstration is credited to Hume. Where philosophers, such as Pritchard, have wished to maintain that evaluation is objective, they have generally abandoned inference from facts and replaced it with the nearest equivalent, a species of non-sensuous intuition which treats moral judgements as if they were factual judgements in their own right.

Reference to non-sensuous intuitions and the attempt they indicate to by-pass rather than bridge the supposed gap between facts and values, leads me directly to the work of G.E.Moore. For it is Moore's argument against what he calls the Naturalistic Fallacy which, more than any other argument, has maintained the importance of the fact-value distinction in Twentieth Century moral philosophy. Unlike Pritchard, Moore does not refer to Hume's argument in support of his own views, but his argument is often taken as supporting precisely the same view that Hume's argument is thought to present — the idea that questions of value cannot be decisively settled.* The assumption that Moore's

* for example, the entry for 'Naturalistic Fallacy' in 'A Dictionary of Philosophy' (ed. Speake, Pan Books 1979) refers to Moore as having named the fallacy, but then quotes the passage from Hume's 'Treatise' in full, as "a much better...statement".

argument is really a restatement of Hume's, or that either argument is a proof of the subjectivity of moral judgement, is quite erroneous. Nevertheless Moore's argument, as it is commonly interpreted, does fall foul of one of the major misconceptions I have already discussed; and his account of the nature of goodness is also influenced by the tendency to treat factual judgements and what resemble them as the sole bearers of objectivity. I shall deal with the latter point first.

In 'Principia Ethica' Moore is really concerned to draw a distinction between two questions.* The first concerns the nature of moral value itself, and the second concerns what things have this value. He thinks that the conflation of these questions has been a major source of error in moral philosophy, because philosophers have failed to distinguish those ethical judgements which are susceptible to proof from those which are not. In his first chapter, Moore argues that the primary moral term 'good' is indefinable, and any attempt to define it (including those attempts to define it in terms of facts about the world, which he attributes in particular to Mill) is a fallacy. This is the Naturalistic Fallacy; but the name is unfortunate, because Moore regards it as a fallacy to identify 'good' with anything whatsoever, natural or not. It is important to recognise that the Naturalistic Fallacy is a fallacy of definition, rather than a fallacy involving any particular kind of definition. Nevertheless, it is the denial that 'good' can be defined in terms of empirical facts which has inevitably most impressed other philosophers.

We might wonder why Moore lays such stress on definition. This becomes clearer when we realise two things. First, Moore conceived of goodness as a property — mainly, I think, because he believed that fundamental statements about goodness could amount to knowledge, and knowledge is perhaps most easily thought of in the form of statements predicating a property of a subject. The decision to treat 'good' as the name of a property was doubtless assisted by the grammatical similarity between 'X is good' and 'X is yellow'. Although predication can be regarded

* (C.U.P. 1903)

as a purely logical relationship between an attribute and what possesses it, it was easy and natural for both Moore and his later critics to treat 'good' as if it really were an analogue of the material-object and physical-property model which is so common when doing philosophy. If 'good' is conceived to be the name of a property then, Moore regarded it as important to show that it was not a compound term for a number of other properties, some or all of which might be natural properties. Second, Moore's insistence on the indefinability of 'good' can be seen as a consequence of his belief that definitions were really forms of analysis, and that 'good' was incapable of analysis, being indivisible. Thus the objection to defining 'good' can be seen as a refusal to identify goodness with anything else at all.

This refusal to identify 'good' with anything else, plus the tendency to treat it as the name of a property, leads Moore into several difficulties. The criticisms of Ethical Intuitionism are well known, and I have already referred to some of them in discussing Toulmin's book. But I must also mention the difficulty inherent in treating goodness as a property which can be known by non-sensuous intuition. In the preface to 'Principia Ethica', Moore denies that he is an Intuitionist "in the ordinary sense of the term"(p.x). By this he means that he denies that anybody knows what things are actually good by intuition. It is only the nature of goodness itself which is intuited. It is well known that other Intuitionists are open to the accusation that when they claim to know that such propositions as 'you ought to pay your debts' are true, and moreover claim to know this by direct non-sensuous intuition, no distinction can be made between the claim that this is knowledge and the suggestion that this is merely a personal belief. (Indeed, if non-sensuous intuition is supposed to be analogous to sensuous intuition, then, as we saw in the last chapter, a good deal of public stage-setting is required before questions of knowledge or belief could even arise.) Because he does not make such claims, Moore escapes this criticism, a feature of his theory which is often overlooked, but he is led directly into another difficulty of a parallel kind. If knowledge

of what goodness is, as opposed to what things are good, is knowledge by direct intuition on a parallel with knowing what yellowness is just by being able to recognise instances of it, there must be a procedure for being able to recognise it aright. I have been at pains in earlier chapters to describe and account for this procedure in relation to colour judgements. But although Moore claims in his preface "I imply nothing whatever as to the manner or origins of our cognition of (intuitions). Still less do I imply...that any proposition is true <u>because</u> we cognise it in a particular way or by the exercise of any particular faculty" (loc. cit.), he <u>does</u> suggest in his Chapter One that there is a straightfoward capacity in people to recognise goodness. "Everyone does in fact understand the question 'Is this good?' When he thinks of it, his state of mind is different from what it would be, were he asked 'Is this pleasant, or desired, or approved?' Whenever he thinks of intrinsic value...he has before his mind the unique object — the unique property of things — which I mean by 'good'." (pp.16-17). This is reminiscent of the classical empiricist's notion of knowledge by direct acquaintance with one's own (private) ideas; and we are quite entitled to ask how Moore <u>knows</u> that that is what everyone, or even Moore himself, has before his mind when he thinks of intrinsic value. Without the possibility of public agreement in judgements, as is the case with colours, we have seen that Moore <u>cannot</u> be sure of this. He can in fact neither know or fail to know what is before his mind, so that his use of the word 'good' to identify it will be devoid of sense. But here again, the assumption that I know from my own case with regard to empirical properties and whatever seems to be just like them is proving irresistable. Thus, in spite of Moore's denial that he was an Intuitionist, the influence of the model of perceptual judgement as the archetype of knowledge was a strong influence on his theory; and by making his notion of goodness a matter of knowledge by intuition, Moore is open to the same criticisms which are decisive against other Ethical Intuitionists.

However, while Ethical Intuitionism has been generally discredited, Moore's "celebrated refutation of naturalism",

as Hare has called it,* has retained many advocates. Hare brackets Moore's argument with Hume's, claiming that behind them both lies the same logical point, that no action-guiding conclusion can be derived from premisses which do not themselves contain an action-guiding element.** (Hare's point is, so far as it goes, true; though I have difficulty in seeing how he found it in Moore's argument. What is not true is that facts never contain action-guiding elements, for facts about my desires do, for instance. It is a fact that I am trying to stay thin, and a fact that cream cakes will make me fat. An action-guiding conclusion about my not eating cream cakes follows with dismal certainty. Hare would be doubly obscure, then, if he were implying here that Moore's argument in some way amounted to the claim that an action-guiding conclusion could never be rationally derived from factual premisses.) But whatever Hare has claimed to discern in Moore, Moore certainly thought he was proving something quite different. And if Moore's argument has anything in common with Hume's, it is a capacity for subsequent commentators to find in it what they want to find. Moore was concerned, as we have seen, to establish the indefinability of 'good'. But his argument has often been taken as showing, not that 'good' cannot be defined, but that moral arguments are inherently unprovable, and hence that the possibility of moral objectivity cannot arise.

Moore's argument for the indefinability of 'good', against the Naturalists' claims to have defined it, falls into two parts, which are best kept distinct. The first part aims to negate any claim to have actually defined 'good', and the second part, sometimes called the open-question argument, suggests a reason why any definition must fail. In the first part, Moore argues that if to call something good meant the same as to call it X (where X is any natural property) then attempts to say that something was good because it was X would result in tautology. For instance, if 'good' meant 'conducive to happiness', then to say that something which is conducive to happiness is good, is to say only that something that is conducive to happiness is conducive to happiness. This is certainly not what Util-

* R.M.Hare 'The Language of Morals' (O.U.P. 1952) p.30
**op.cit. pp.28-31

itarians such as Mill were trying to say when they argued that all and only what promotes the general happiness is morally good. But any attempt to link 'good' to a natural property, or set of properties, by a definitional relationship must break down in this way. Whether Mill's notorious argument is after all an attempt to define 'good' in terms of human desires is a matter of dispute. But if it is, Moore's argument is effective against it.

The second part of Moore's argument is supposed to show why 'good' cannot be defined. This, for him, is because 'good' is the name of a property which is indivisible and sui generis. Because moral philosophers no longer treat 'good' as the name of a property, however, this aspect of the argument is now usually ignored, and its importance is taken to lie elsewhere. It is felt to lie in his remark that, of any attempt to define 'good', "it may always be asked, with significance, of the complex so defined, whether it is itself good."* This is of course the converse of the previous argument, and again it really concerns definition. 'Good' can be shown in advance not to be equivalent in meaning to any set of natural predicates; and so any attempt to define 'good' must fail. From this it follows that nothing in the world is good just as a consequence of the meaning of 'good'. For everything in the world consists of natural properties only, and 'good' cannot mean the possession of any of them. Yet though Moore's argument shows that nothing with only natural properties could be good by definition — a logical point — his argument is often treated as saying something quite different, and as making an epistemological point. By showing that any statement of the form 'X is good' will be synthetic, not analytic, Moore has been taken as implying that any judgement of this form is open to question (which it is, but no more so than any empirical judgement), and that any argument with factual (natural) premisses and an evaluative conclusion must be open to doubt or actually invalid. This shift in the reading of the argument is due at least in part to the fact that we no longer regard 'X is good' as a statement.

*'Principia Ethica' p.15

Thus it would now be thought inappropriate to call it either synthetic (as Moore does) or analytic. But that does not change Moore's essential point, which is that an evaluative judgement cannot be validated by an appeal to a definition, so that in consequence such a judgement is always open to the demand for a justification to be given in some other way. That any evaluative claim, however we arrive at it, is open to doubt in this way cannot be disputed. But I have already argued that it is a mistake to suggest that this has any special significance in the case of evaluative claims, and seriously misleading if it is implied that evaluative claims would have to meet the requirement that they should follow by definition from some fact. It would be absurd to make it a requirement of an argument with factual premisses and an evaluative conclusion that the conclusion must follow by definition from the premisses, seeing that this is not a requirement when we argue from factual premisses to a factual conclusion. If we lose sight of the original aim of introducing the open-question argument, which was to show that nothing about what is, or is not, good in the world can be determined from considering what 'good' means, we may also lose sight of the fact that the question 'but is it good?' "may always be asked with significance" because no judgement that some state of affairs in the world is good can be justified by reference to a definition, and not because the judgement is an especially doubtful or unprovable one. Moore himself believed that it is possible to say what things are in fact good, and indeed he makes claims of this sort in 'Principia Ethica'. He would hardly have done this, no matter what justification he thought could be offered for such claims, if he had already argued in the first pages of the book that no such claim could possibly have any justification.

Of course, the belief that the facts about the world by themselves, independent of anybody's attitude to them (or, as Moore would have it, independent of some form of non-sensuous cognition relating to them), cannot give us a reason for making an evaluation, is not incompatible with Moore's position. Nor is it incompatible with my own. I am not trying to suggest that reason alone can enable us to

argue from facts to values. Instead, I am resisting the ideas that (a) only if it could do this would moral objectivity be conceivable, and (b) the fact-value distinction shows that moral objectivity _is_ impossible by, in this case, suggesting that any moral conclusion is open to a specially powerful form of doubt, and is really unprovable in consequence. Part of the reason for this widespread belief in the inherent unprovability of moral conclusions stems from misunderstanding the significance of Moore's open question in the way I have suggested. For it _looks_ as if the question might be applicable to _any_ reason given in support of a moral claim, and so suggests that no moral argument could be conclusive. But insofar as the open-question argument, and with it the whole Naturalistic Fallacy argument, comes to be used as an argument against the possibility of there being proof for moral beliefs (and hence moral objectivity), it is open to just the same objections that applied to the original argument for the fact-value distinction, earlier in this chapter. There we saw that while the synthetic character of factual judgements made disagreement a logical possibility, this possibility was never held to preclude decisive conclusions on what was, or was not, really the case. However, the identical logical feature of evaluative judgements _was_ held to preclude decisive conclusions on what was, or was not, good, since it became covertly transformed into an epistemic principle of systematic doubt. And the same illicit procedure is being followed when the open-question argument is referred to as a proof that value-questions are always open _in the sense of being objectively undecidable_.

To illustrate this point, and perhaps make it clearer, let us compare two inferences, one factual and the other moral, in order to see what impact the open-question argument has on them. The moral example must be somewhat circumscribed since, as I suggested earlier, moral judgements tend to be inherently more complicated than factual ones. Suppose someone says 'if you put a yellow object under a sodium-vapour lamp, it will look white', and I ask myself whether this is so. The obvious move is to try the experiment and see for myself. This is because how the

object looks to me is normally decisive, since I am a standard judge. When I see that the object does look white under those conditions, it is still open to me to ask 'but does it, after all, really look white?'. That is to say, since the judgement is a synthetic one, it is not self-contradictory to repeat the question. But we must note that, unless I suspect that the observation I made was in some way abnormal, there is no reason to ask the question again. If I do ask it again, it must be in the context of exploring or explaining a possible abnormality. The only other reason for my asking the question repeatedly is that I simply do not understand what counts as proof in cases of this sort.

Now let us imagine that, like the Good Samaritan, I find an injured man at the roadside, and the idea comes to me that if I tend his injuries I will be doing good. Here I might also ask the open question, for it is not self-contradictory to doubt whether doing this will be to do good. (We must notice that there are a number of other questions I might also raise at this point. I might wonder whether he really is injured, or really suffering on account of his injuries. I might also wonder if I ought to help him (because I have other pressing obligations at this time), or if I can help him, or if it would be prudent to help him, or even, in some conceivable cases, whether it would be just to help him. But these are all different questions from the one we are considering.) Now of course I do not deny that we can, without self-contradiction, doubt this moral conclusion that helping the injured man will be a good act; nor do I deny that it will not be self-contradictory to ask the open question no matter what reasons have been given in support of my moral belief. But it does not follow from this that the question cannot be decisively answered, just as it does not follow that the question previously raised about the colour of an object cannot be decisively answered. If, for example, the relief of human suffering is one of those cases where if anything is a case of doing what is in fact good, then this is, it will be pointless to go on asking whether to tend this man's injuries will be, after all, a case of doing good. To go on asking the question

here will then indicate a lack of understanding of when a further request for justification is appropriate and when it is not, in the same way that the man staring at the book which he can see looks white, and repeatedly asking himself 'but does it really look white?' shows a lack of understanding. For if anything is a case of looking white, then this is.

It will be objected that there is a different, and very important, sense in which someone might significantly doubt whether helping the injured really is a good thing to do. I am not thinking at the moment about doubts as to the consequences of helping the injured, or doubts about what would happen if everybody did this. For these considerations are part and parcel of what is involved in considering the action's goodness; and so to settle the question of whether the action is a good one decisively is to settle these questions decisively too. Rather I am thinking of the genuine possibility that someone might wonder whether helping the injured was a principle that belongs to a false moral creed (a 'slave morality', perhaps). In this case, it will be insufficient to reply that someone who wonders about this simply does not understand what it is to do good, and so cannot see that his doubts are pointless. It may be true that his doubts point to a lack of understanding on his part; but can we always be sure of this?

There is no short or simple answer to this objection; but there <u>is</u> an answer, which I hope will emerge more fully in the course of the next two chapters, and which can be at least indicated here. In the first place, it is significant that doubts of this sort are characteristically expressed as part of a different moral creed, rather than felt by someone who shares our own moral code but feels a doubt in just this case. We would think it odd, for instance, if a man showed concern for the suffering of his fellow men in every way except this one — he is concerned to prevent or ameliorate fatigue, for example, because of the suffering involved, but is utterly indifferent to that caused by physical injury. We would think his attitude abnormal, and wonder whether he really understood what physical injury is like. But on the other hand, a moral creed

which shows a systematic indifference to human suffering, or to certain broad categories of suffering, is at least consistent. We do not wonder here whether someone who holds such a view understands what injury is, but whether he understands what moral goodness is. If there are no points of similarity between his beliefs and ours, then I maintain that we are entitled to accuse him of a lack of moral understanding, rather than just saying that he has a different moral attitude to suffering. But if we find his moral beliefs to be very like our own in other respects — if he places great importance on keeping promises, for example, and acknowledges rights of property as we do — then we must recognise that this man does have a sense of what is moral, albeit, we may think, a limited and partial one; so we cannot simply say that he lacks any understanding of what is constitutive of morality. But if that is the case, then there must be many other situations where he, like us, would say that there is no further scope for doubt about whether acting in a certain way is, in fact, good. He would agree with us, for instance, that keeping a promise is a good thing to do, and that once one has seen that doing X is a case of keeping a promise, the question of whether doing it is, after all, good has been decisively settled. So here we can see that, in admitting that what he has is a morality, we are recognising that he would give decisive answers about what it is good to do in a whole range of circumstances where we too would give decisive answers. (I shall argue later that these answers will also be the same answers.) This being so, we can see that the open-question argument does not allow anyone who has an understanding of what moral goodness is just to continually express doubts at every juncture about what things are good. There may be areas where others would express doubts but we would not, and vice versa. But what is impressive is that there must be a constancy of moral response to a wide range of actions and situations, and not the fact that there may sometimes be a variety of response. So my suggestion is that the man who really does doubt whether it is good to relieve human suffering is responding in an abnormal way by doubting; and if this is not an abnormal doubt on his part, it is because he

lacks a full understanding, or perhaps any understanding at all, of what moral goodness is. Either way, this example does not show that the open-question argument is a systematic principle of moral doubt, so that the notion of there being justified moral beliefs is illusory.

In suggesting that there are decisive answers to a good many questions about whether some action or state of affairs constitutes a prima facie moral good, I am raising directly the crucial issue of why some considerations should count as relevant and decisive in answering moral questions. The recognition that there can be decisive evidence for a moral conclusion does not, as we shall see, fully explain the basis for saying that moral judgements can be objective. For there is the dimension of moral meaning also to be considered. Before leaving this discussion of the misconceptions that surround the fact-value distinction, therefore, I want to turn to one last example of the distinction which, apart from incorporating some of the errors I have already discussed in this chapter, does focus our attention on the questions of moral evidence and the meaning of moral terms. (This last point is the other misconception that I promised earlier on that I would refer to.) The argument is put forward by Toulmin, in that chapter of his I have already referred to, and draws a contrast between factual and moral judgements in terms of the possibilities for disagreement which are open to two men making such judgements. The central passage runs as follows:

"If I am confident that both men are candid and in full possession of their faculties, and that they employ the same language, dialect and usage (i.e. if all the sources of disagreement over simple qualities are removed), there will be no point in my asking whether they agree or disagree about the colour of a pillar-box: there is no room for disagreement. If, in addition, I know that they have counted together the sides of a given polygon, it will be as pointless to ask whether they agree about its 259-sidedness. But, though I know all this, it will still not be silly to wonder, for example, whether they will agree that meekness is good, or that such-and-such is the right de-

cision. Even if there is neither deception nor defect on either side, even if both parties are fully informed about the case and both mean the same by 'good' and 'right', it still makes sense to enquire whether their moral judgements are in fact the same.

This difference between values and properties is crucial..."*

Here we have what appears to be a carefully drawn contrast between factual and evaluative judgements. Also, as we might now expect, there is present the element of vagueness about the status of the possibility of disagreement between the two men in their moral judgements. As I have argued, it <u>always</u> "makes sense to enquire" whether two people are in disagreement, whichever type of discourse they are engaged in, because that only shows that their judgements are not analytic. But Toulmin is muddling this logical feature of judgements of both the types we are considering with the quite different idea that, when certain conditions for judgement are satisfied, we have adequate grounds for judging with justification, and it is in <u>that</u> sense that there is "no room for disagreement". This confusion allows Toulmin to imply that if two people disagree over a moral question, it does not follow that at least one of them is in error. Now this, I have argued, is an epistemological claim of some importance, and cannot just be asserted. To prove his claim, Toulmin would have to show that in moral judgement there can be no such thing as being in error. He has not, as he thinks, shown this merely by pointing out that disagreement is logically possible here.

Turning now to another aspect of his argument, it is interesting to notice that although Toulmin is quite careful to describe his examples of factual judgement thoroughly so that there is "no room for disagreement", he does not in fact draw attention to one of the principal reasons for this being so. It is not just that he has carefully excluded all the reasons why an empirical judgement can go wrong. The reason why the factual judgements leave no

* Toulmin, op.cit. p.20

room for disagreement is that there are decision procedures at work there which are universally accepted as the relevant ones. Looking is normally decisive when an individual judges the colour of something; careful counting is similarly decisive when it is a question of how many sides a polygon has. These procedures are decisive because they are obviously the relevant ones, and because "all the sources of disagreement" have also been removed. But when, on the other hand, we turn to the question of whether meekness is good, the situation is not at all a comparable one. Toulmin has attempted once more to remove all the extraneous sources of disagreement (but see below); yet this time there is no mention or suggestion of there being any relevant methods of deciding the question. It is not just that these methods are not mentioned; their absence is taken as much for granted as their presence was when the judgements in question were factual ones. And this is hardly surprising when we also notice that, whereas the examples of factual judgements were highly specific, so that it was <u>perfectly obvious</u> what the relevant method of decision should be, of the two moral judgements which Toulmin proffers as examples, the second is <u>entirely unspecific</u>, and the first refers to the goodness of a virtue considered as a general proposition. Toulmin's argument would have looked much less convincing had he also given examples of equally general empirical judgements. (Questions like 'are the seas blue?' come to mind as rough empirical analogues of 'is meekness good?'.) He tells us that "it will not be silly to wonder" whether there will be agreement on these questions — and indeed it will not, since one of the conditions necessary for answering them is rendered so remote as to be almost inconceivable. What we are left wondering about is how to make these questions precise enough for the relevant methods of decision to get a grip on them at all. (As a first step at least, we would want to consider what our response would be to specific cases where the question 'should I act meekly here?' arises.) Faced with this degree of vagueness, it is easy to surrender to the idea that the sort of moral disagreement Toulmin envisages is simply irreconcilable, instead of just being a logically possible feature of all non-analytic judgements.

So much for Toulmin's failure to see the significance of the decision procedures he mentions, and his failure to refer to those he leaves out. But there is a second point about his argument which also deserves closer examination. He includes in his list of "all the sources of disagreement" for moral judgements the inevitable, and inevitably misleading, reference to both the parties being "fully informed about the case" (so hinting that in a <u>really</u> decidable dispute over moral judgements, it would have to be the facts themselves which were decisive, because only facts <u>could</u> be). But he also mentions the idea that "both men mean the same by 'good' and 'right'". It transpires from the next paragraph that all he means is that there should be no difference in the degree of approbation or disapprobation intended by the use of a particular word, so that we are not misled when somebody calls something 'not too awful', but says it in an ironic way, really meaning 'good'. This, as Toulmin admits, is trivial. He is just being tidy about the possible sources of disagreement as he envisages them. There is, however, a different and decidedly non-trivial way in which the question of agreement or disagreement about the meaning of moral terms can arise. This involves the idea, which we arrived at in the previous chapter, that to understand what a word means is to be able not only to define it but also to use it when and only when it is appropriate to do so; that is, when, other things being equal, the judgement in which it occurs will actually be correct. Now the sort of meaning Toulmin has in mind when he speaks of two men agreeing on the meaning of 'good' is restricted to the definitional level, where all that is specified is the degree of approbation or disapprobation intended. Here, then, the most that is required is that both men should recognise that 'good' (in moral cases) is a term for a certain degree of approbation which can be applied to anything that could count as a moral case. But could 'morally good' mean only this and still be the sort of term where it makes sense to say there could be agreement or disagreement? We can only talk about agreeing or disagreeing over whether a pillar-box is red or a polygon 259-sided as long as we

recognise that something counts as finding out, and we agree on what that something is. Similarly, we can only talk of agreeing or disagreeing about whether something is morally good if there is something built into our notion of moral goodness which gives us the idea of what would count as a right answer and what would not, at least in ordinary cases. If this is correct, Toulmin is wrong to suggest, as he does, that two people could agree on the meaning of 'good' and yet regularly disagree about what things are in fact good without at least one of them being in error, or else showing a less than adequate understanding of what it means to call something good.

There is, then, a problem about the meaning of moral terms encapsulated in Toulmin's argument, although I am sure Toulmin does not realise this. It is a problem intimately related to the question of there being relevant and decisive evidence for moral conclusions, which I have already touched on. I shall turn to a discussion of some questions about evidence in the next chapter, and this will take us a considerable way towards understanding the place of objectivity in moral judgement. It will not take us all the way, however. For we will also need to see how it is that the decisiveness of evidence stems from the requirements imposed on our judgements if the moral terms we use are to be meaningful.

These, however, are topics for the next two chapters. I have introduced them here because they arise naturally from the misconceptions inherent in the arguments I have been considering in this chapter. The belief which these arguments express, that moral judgements are so different from empirical judgements that they cannot be objective, has tended to obscure the issues of moral evidence and moral meaning which are central to an understanding of moral objectivity (and <u>mutatis mutandis</u> to evaluative objectivity as a whole). In this chapter I have tried to bring out what I see as the two main faults within the fact-value distinction, as it is commonly drawn — the confusion of epistemological and logical questions about the openness of judgements, and the assumption that the objectivity of moral conclusions could have its source only

in factual premisses from which the moral judgement must be deduced. Both these faults, and other less drastic ones associated with them, may be found in arguments of various degrees of crudity or sophistication. But whether subtly or naively presented, the underlying reason for this entire way of seeing things is, as I suggested at the very start of this thesis, the belief that being objective is, in the last analysis, really just the same as being factual.

CHAPTER SEVEN

EVALUATION, EVIDENCE AND REASONS

Towards the end of the last chapter I referred to two important possibilities which suggest themselves to us when the misconceptions surrounding the fact-value distinction are cleared away. The first of these possibilities is that there might be facts which constitute decisive grounds, or amount to conclusive evidence, for some moral judgements. The second possibility is that our ability to use moral terms intelligibly presupposes a significant measure of agreement in the moral conclusions we arrive at. Both these possibilities must be explored, since both point to the existence of at least an objective dimension to moral discourse; and the remainder of this thesis will be an attempt, if not to chart this dimension, then at least to give some indication of its nature and extent in the course of arguing a case for its existence.

Although both the possibilities I have mentioned point to the existence of objective moral judgements, they do so in different ways and to differing extents. The first of them, which is the subject of this chapter, will lead us to see that agreement in our basic moral judgements is necessary if moral discourse is to have the rational characteristics that it really does have. However, this is not by itself enough to show that moral judgements can be objective, although it is enough to show them to be interpersonal, and the necessity of interpersonal agreement in judgements is a very strong indication of objectivity. However, as we have seen in earlier chapters, objectivity requires more than that there should be agreement in judgements, however widespread. It requires also that this agreement should be justified agreement; and the basis for justified agreement is not to be found in the fact of agreement but in the explanation of why this agreement is a normal feature of judgements of this type. Thus this chapter provides proof of the existence of moral objectivity only when taken in conjunction with the arguments in the next one.

I have spoken so far of there being some moral judgements, or else some types of moral judgement, which are objective, and it is time to make what I mean by this a little clearer. I am not, of course, suggesting that every moral judgement is an objective one. This would be an absurd claim, just like the claim that every empirical judgement is objective, because it ignores the possibility of human error, idiosyncracy or lack of understanding. It also leaves out the important point that we often have to make a **judgement** without being in full possession of the facts necessary to make it properly. Indeed, the case has sometimes been made against the possibility of moral objectivity on this latter ground alone. Bernard Williams has argued, for instance, that a man's moral decision might turn out to be right or wrong as his future luck dictates.* He cites Gauguin's (presumably) moral decision to abandon his family for the sake of Art. Williams says that the painter's decision could only be justified on the grounds that he did, as it happened, turn out to be a great painter. But since Gauguin could not have been sure in advance that this would be the case, it is impossible that his decision could have been an objective one. The example does not prove all that Williams intends, however. He would only have disproved the possibility of moral objectivity if he could show (a) that all moral cases were like this with regard to the facts, and (b) that the moral rightness of an action is always dependent on its consequences, and never on any other consideration. Leaving the second point aside, we can see that the Gauguin case misses the central difficulty about claims to objectivity in such cases. For sometimes all the facts about a moral dilemma are in (as they now are with respect to Gauguin's desertion of his family), but the question of whether he was right to act as he did remains. A case of this sort is still difficult when all the facts are known because it is essentially a problem of deciding between two courses of action which are incompatible but

*'Moral Luck' in 'Philosophical Papers 1973-80' (C.U.P.)

which both have, or lead to something that has, moral value in its own right. And it is not this sort of case that I want to say has an objective solution. For this sort of dispute characteristically arises from a clash of acknowledged but conflicting values; and while I do not want to go so far as to deny that such dilemmas ever admit of objective solutions, it is a contingent matter that they do so. I shall discuss this point in the final chapter. Here we need only to realise that the type of case Williams cites is very different from, and actually presupposes, the sort of moral judgement which I wish to claim can be objective. For the Gauguin case is only a moral dilemma at all because it is a dilemma between two acknowledged goods or obligations. If there were not something to be said in moral terms for each of Gauguin's possible courses of action, his situation would not be a moral dilemma at all, but a psychological conflict between duty and self-interest. This would raise interesting questions, but they would not be questions about moral objectivity.

The type of judgement for which I want to claim objectivity, then, is the kind of judgement which has already been made before Gauguin's dilemma can arise as a moral dilemma. It is the judgement that a particular state of affairs or course of action has moral value in its own right. Thus in Williams' example it is presupposed that possession of creative talent leads to something good in its own right and so imposes an obligation on its possesser, and being a husband and father imposes a responsibility on a man for the good of his family. It is judgements that these things are so which I shall argue are objective. It will be noticed that objectivity does not entail indefeasibility, and it is for that reason that I am concerned with goodness rather than rightness or obligation. In some cases these last two terms do imply that a decision between competing values has been made, whereas 'good' does not. It could be said that my concern is with prima facie value rather than with relative value, a distinction which has been rather lost sight of in contemporary moral philosophy.

I am not arguing, then, that to every genuinely moral question there is a right answer, even in principle. I am only claiming that ascriptions of *prima facie* moral goodness are objective; and these, together with the mechanisms of reason-giving and reliance on evidence which they give rise to, provide an objective background against which productive and meaningful moral discussion can take place.

It is often said that there is certainly one area in which there can be decisive evidence for evaluative judgements. This is the area where judgements relate to instrumental or functional things, processes or activities. To say that a knife is sharp, for instance, is certainly relevant to, and in some instances decisive for, calling it a good knife. To say that a farmer's crops and herds flourish as a result of his husbandry is, except in special and parasitic cases, reason enough to call him a good farmer. And if a sprinter holds the world record and a string of Olympic medals, we can hardly deny that he is a good sprinter, since these achievements mark him as the best (or one of the best) sprinters of his day. Criteria for the proper evaluation of cases such as these might be supposed to derive from the meaning of the term for whatever is being evaluated. Thus a knife, being an instrument for cutting, and being this as a matter of definition, can only be evaluated *qua* knife in respect of its cutting powers (though it must be added that a full evaluation of a knife *qua* knife will refer to such things as lead to cutting consistently, safely and accurately as well as to the primary criterion of sharpness). Similarly, a farmer or a sprinter can only be evaluated as these things by reference to the criteria for doing these things well, properly or efficiently; and these criteria are determined by what it is to *be* a farmer or a sprinter.

It is, however, a mistake to regard criteria for evaluation as being a consequence of *definitions*, except in those special cases where objects are defined solely by their function. For example, we could not expect the definition of 'man' to yield moral or even functional

criteria for human goodness. Possibly a fairly detailed description of man might indicate how he functions physically, mentally and socially. However, since man does not have functions in the sense of being for something, or having some particular purpose in being just a man, even this description will not yield criteria for being a good man in a functional sense (unless we regard man's function as being to survive); and this process certainly will not yield any moral criteria for man. In 'Goodness and Choice' Mrs Foot moves from consideration of cases like 'knife', 'eye' and 'lung', which are functional in a strong sense, to examples which, she claims, "show that the range of words whose meaning determines criteria of goodness is much wider than that of functional words".* But Mrs Foot is using the word 'meaning' in rather a broad sense. 'Father', for instance, is not usually defined as a person who has a concern for the upbringing of his children, although we might accept that this concern is part and parcel of being a father, and so yields criteria for being a good father in that way, viz. by displaying this concern in appropriate ways. Mrs Foot accepts that what counts as an appropriate display of concern will vary from society to society, and that opinions about it may differ, but insists that there are limits to what could count as appropriate forms of concern. Ritual sacrifice of one's children might be a sign of piety, for instance, but could not count as grounds for calling one a good father. She does not say so, but clearly has in mind that only what will further the welfare and development of the child could count as an appropriate form of concern. She does refer to "caring", which I take it embraces those concerns. Now we might perhaps be able to invent circumstances in which, say, ritual sacrifice of one's children was a sign of caring for them, though this would involve the father in some rather unusual beliefs. But Mrs Foot's point is preserved here, since we can say that, given the beliefs he had, the father was still acting in what he believed to be his children's interests, was still caring for them, and was in that sense at least trying to be a good father. So it is still only with

* PAS 1961, Reprinted in her 'Virtues and Vices', p.137

regard to whatever counts as caring that a man can be a good father. This has taken us quite a long way from the meaning of 'father'; but given that we understand the connection between being a father and whatever counts as caring for one's children, the word 'father' will yield criteria for goodness in the way Mrs Foot describes.

Now if this is the sort of case which allows us to find determined criteria for making evaluations which begin to look like moral evaluations, we can see that what matters is not meaning in the sense provided by a definition, but meaning in a sense which implies an understanding of the thing or activity in question as having a particular point. Mrs Foot emphasises this in one example in which she speaks of the criteria for good coal as something "that depends on what coal happens to be used for" (op. cit. p.138). In a society where coal was mined for other purposes than as fuel, the criteria for its goodness would be different from ours, "while outside such a context it is not clear how anyone could talk about coal as good or bad at all". (It is not clear that anyone would want to talk about 'coal' having a different meaning in these two societies, incidentally; and elsewhere meaning disappears from Mrs Foot's discussions as a source of criteria for goodness.)

For the most part, then, Mrs Foot's notable defence of Descriptivism does not rest on strict considerations of what words mean. It is our knowing the point of a particular thing or activity, or our knowing the point of being something, or our having a particular interest in something, which is the real source of determined criteria for goodness, especially in those cases such as being a good father, where we are inclined to say that the sort of interest we have makes the evaluation one that we would ordinarily class as moral.

In 'Moral Beliefs', a rather earlier article, Mrs Foot argues further that seeing the point of a thing or activity, or having a particular interest in it, can be seen as having a connection with what a man wants.*

*PAS LIX 1958

So moral judgements can, she thinks, be shown to be determined by recognising, first, the fact that criteria for evaluation are determined by their connection with our seeing the point of, or having an interest in, various things and activities, and second, recognising that there are things and activities that all men want or have an interest in.

The second stage of this argument has been much attacked. It has been suggested that statements about men's common interests will turn out to be too vague to be informative. Again, Phillips and Mounce* argue that the sort of moral decisions men actually make shows that there is nothing, not even freedom from physical injury, which all men can be said to want; and that in any case what people want is often determined by their moral beliefs, and not vice versa. Their first argument does not seem to be conclusive. All men might indeed want freedom from physical injury, but might on occasions want other things more, and accept physical injury as part of the price that has to be paid for what they want. However, the second part of their argument is a powerful objection to this sort of Descriptivist ethics. A different objection, but also a powerful one, is that an Interest theory of this sort will always yield prudential rather than moral reasons for doing anything, so that it cannot be a <u>moral</u> theory at all. If a person acts in a certain way because it is in his interest so to act, or because by so acting he will satisfy his wants, it is hard to see how he could be said to be acting disinterestedly; and this is a mark of a moral act, since to be moral one must sometimes act against one's own interests. Mrs Foot might reply that by acting against one's own particular interests, one is acting for the general interest, which is also one's own. However, either we are left here with a conflict of interests which, if irreconcilable, destroys the idea that morality can simply be equated with interest; or else a man can never really act against his own interests, which seems to undermine the idea that to act morally is sometimes to act in a different way from what one would do if one were motivated only by

* 'On Morality's Having a Point' (Philosophy, 1965)

what one wanted for oneself.

The second part of Mrs Foot's theory, the notion that criteria for moral judgements are determined because there are things that all men want or have an interest in, looks distinctly insecure, then. But we do not need to accept the second part of the theory in order to accept the first. We can accept that if people have an interest in something, that will determine their criteria for evaluations of whatever is relevant to their interest. But we do not need to accept that people all have the same interests, or that human interests are necessarily limited to what people want for themselves. We recognise, for example, that if a man wants power above all else, then whatever there is in the world that is a means to power will be valued by him, and by anyone who shares his interests. And we can understand this without having to share those interests ourselves. This fact is to some extent obscured by the way in which Mrs Foot uses the term 'interest' in both halves of her theory.

R.M.Hare, however, Mrs Foot's principal target, has argued against both parts of her theory; and his argument is in one way quite correct. In 'Descriptivism'* he points out that to describe something (such as a man's actions in battle) in a certain way may well be grounds for commending him (e.g. as courageous). But that is only because "we already have the standards of values according to which to do that sort of thing is to display outstanding merit". Our standards could be other than they are. And if a particular choice of descriptive terminology seems to impel us to pass one judgement rather than another on what we are describing, we can always describe it using other words which do not impel us in this way.

Even though this suggestion of Hare's might not always be very easy to put into practice, it does cast very serious doubt on the idea that, even if we can describe some action or event in entirely neutral language, we will still all be impelled by what has been described to view it from one

* Proceedings of the British Academy, 1963

particular standpoint, as having a particular point or significance for all of us, so that the neutral description itself will be what determines our moral response. But that is what strict Descriptivism would require. Yet no fact by itself, and apart from the interest we have in it, can determine our evaluative response. It could always be the case that you and I would see the same facts as having different significance, or that others would respond differently from either of us. And this is not quite the empty logical 'could' discussed in relation to the Open Question Argument in the last chapter, because in view of the claim Descriptivism would have to be making about the relation between any correct description and one particular evaluative response to it, the question of demonstrable error is raised. According to the first part of Mrs Foot's theory, a man who derives his evaluative criterion from some fact should be able to point out where someone who doesn't do this has gone wrong. But this is only possible if the other person shares his interests; and that need not be the case if the second part of Mrs Foot's theory is incorrect, as I have argued it is.

In the light of all these difficulties, what is there left of Mrs Foot's analysis of the relation between description and evaluation? I suggest that her notion of there being an internal relation between the meaning a descriptive term has and the kind of valuation which someone who uses that term must put on whatever he describes by it, is in itself a valid one, providing that we take 'meaning' in a very broad sense which includes an awareness of the point or interest that the thing meant has for us. Mrs Foot's mistake is in supposing that this is the case while at the same time regarding the description as neutral in itself. If things were like that, every correct description of an object, process or activity which has point, purpose or interest for us would yield criteria for evaluation which would necessarily be the same for everyone, and all value-judgements would then be objective providing no errors of reasoning were made, and providing also that everyone was clear about the point or interest involved. In these circumstances,

the objectivity of evaluation as a fact about the nature of human judgement would scarcely be a topic of philosophical dispute in the way it has been. Now as it happens, descriptions of the functional objects from which Mrs Foot's discussion starts do appear to be neutral descriptions in the required sense, and so appear to provide a model for other kinds of description which give rise to value-judgements. Let us return to the simple case of what makes a knife a good one. There are two points to be noticed. First, because knives are functional objects, the interest that we have in anything we call a knife is already fixed for us. That is why the definition of 'knife' yields criteria of goodness so readily. For both definition and criteria are consequences of our interest in having and using knives. Second, and pace Mrs Foot's remark to the contrary, our notion of what counts as cutting well, and consequently our idea that sharpness is the essential criterion for goodness in knives, is not a fixed one. True, the individual cannot rationally deny that good knives should be sharp, and still claim that knives are for cutting and that a good knife should cut well. For 'knife' and 'cutting' and 'cutting well' are terms in a language he shares with others, and these terms have an agreed application for users of that language. In that sense, Mrs Foot is right to say that "no-one is at liberty to pick on just any kind of cutting as cutting well".* But a community or a culture is, in a certain sense, at liberty to decide what its interests are, what means it will employ to further those interests, and what, for it, is to count as having succeeded in furthering those interests with the means it has adopted. No doubt, given that we are the creatures we are and that we all inhabit the same physical world, notions such as 'cutting with a knife' are liable to be much the same everywhere. But that should not blind us to the fact that what underlies our acceptance of sharpness as the criterion of goodness in knives is the interest we all share in doing what the objects we call knives enable us to do, and the agreement we all share in what counts as furthering or succeeding in

* Goodness and Choice, op.cit. p.135

doing what we do with knives. Both sorts of agreement underlie the fact that criteria of value in knives are determined.

Now if agreement in aim or interest, and also in what counts as furthering it, are both necessary if the evaluation of functional objects and the like is to be determinate, this will be no less true of other sorts of evaluation, such as the making of moral judgements. While we cannot point to anything which might be said to constitute the moral aim or interest of mankind, we can certainly point to a number of very common aims or interests that people actually do have, and which we can recognise as moral aims or interests rather than as prudential ones because they may on occasion run counter to our own personal interests. The prevention of suffering in others is as a matter of fact a fairly common non-prudential aim. Given that this is an aim which many people share, and given also that they agree that a particular action X does prevent human suffering, the question 'is it good to do X?' does have a determined answer in the affirmative for them. Thus 'he does X' would be indisputably evidence for 'he is a good man' amongst those who share the aim under discussion; and 'he does X' would be conclusive evidence for the judgement 'he sometimes does a good deed'.

But this account of the circumstances in which factual evidence for an evaluation can be conclusive presupposes both the sorts of agreement I have mentioned. We cannot just say that such agreement is a fact about human beings, which in effect is what Mrs Foot does when she claims that there are things that all men want. However, the fact that people do regularly make the same evaluative responses to facts of various kinds does suggest that both sorts of agreement do occur. What it does not show is that both kinds of agreement must occur. Now if the necessity of such agreement can be shown to be a requirement if our moral discourse is to function as it does, this will show that some facts can count as conclusive evidence for a moral judgement. For if in order to preserve a feature of our moral discourse which is central to it, there must be such general agreement, this agreement will ensure that

there are criteria which will determine judgements among all of us who share the agreement; and therefore evidence for a judgement must sometimes be decisive — for us. If agreement of these sorts is necessary to moral discourse as we understand it, it follows that moral judgement is an interpersonal matter. This is not enough, as I have said before, to prove that moral judgements can be objective as such. For that it must be shown that our agreement is not a mere consensus of opinions, but can be backed by sufficient justification for us to say that others who judge differently are actually wrong. However, demonstrating this necessity of there being agreement of the requisite sorts is an important step towards proving that moral objectivity does exist.

The necessity for agreement in the aim or point of a thing or activity which we are evaluating, and also the need for agreement on what counts as furthering the aim or point, can be shown by considering the role played by reasons in moral discourse. Moral thought and moral discourse is rational in at least two ways. First, it involves inferences about the consequences of actions, and the relationship between particular acts and general principles of action, in a way that shows that the mental processes involved in moral decision-making are rational, and not just intuitive or appetitive. We would not consider that a man, even a saint, had a fully developed moral faculty if he just did things, but never considered why he did them, what the consequences of his doing them might be, or whether such actions instantiated a moral rule for himself, or for everybody. However good his actions were, however much goodness actually resulted from them, and however inspiring his example might be to others, we would be entitled to wonder if he really knew what he was doing when he acted as he did. Secondly, if our saint not only did not but could not make judgements on a rational basis, he would be unable to engage in a central aspect of the moral process. Moral decisions often involve doubt and reflection, and <u>ex hypothesi</u> he would be in no position

to weigh up alternative courses of action or, just as important, to explain or justify his decisions to himself or to others, or to persuade others to act contrary to their interests when morality demands this. Since his own actions are unreflective, it might be doubted whether those actions even deserve to be called moral actions, although they might seem to others to be moral in that they resemble genuinely moral actions. The saint himself, however, would not be able to distinguish between doing something for moral reasons and doing something just because he felt the urge to do it. This difficulty is inherent in his inability to give reasons for his actions.

Both the need to give reasons, and the sort of reasons that are given, sharply differentiate subjective and objective expressions of attitude. This is not to say that reason-giving plays no part in discourse relating to the individual's personal preferences or likings. It is quite proper to ask someone why he likes something, or why he prefers one thing to another, when the only thing in question is his personal tastes. A man might respond to a request for the reason why he chose strawberries rather than gooseberries by saying that he prefers strawberries on account of their sweetness, whereas gooseberries are so tart. But this is only superficially a reason, for it only indicates more precisely that aspect of the fruit to which his attitude relates. The reason given is informative and to that extent we are justified in calling it a reason. But it does nothing to explain his attitude in the sense of taking us any closer to understanding why sweet things might be preferable to tart things. This is apparent when we consider that the man whose tastes are quite different, and who prefers gooseberries to strawberries, can give as his reason exactly the same features of the taste of the two fruits, the only difference being that <u>he</u> likes fruit to be sharp-tasting. These reasons provide no justification, nothing, that is, that would rationally impel us to adopt one or the other of the two preferences for ourselves. There <u>are</u> no justifications for personal tastes in this sense, and the reason a man has for choosing

one fruit rather than another is simply that he has the preference that he has. His 'reason' characterises that preference but does nothing to support it.

The lack of any need for a justification for saying that one likes something, providing one is sincere and not deceiving oneself, is clearest in connection with extreme or bizarre preferences. If a man sincerely enjoys doing something thoroughly odd, such as hitting himself on the head with a brick, there may be no rational explanation for what he does beyond the fact that he enjoys doing it. And even if some explanation is forthcoming, in terms of what sort of satisfaction it gives him to do this, few other people would, I imagine, be persuaded to do it too, unless it happened that they wanted that sort of satisfaction too, but had not realised that here was a way of obtaining it. We might well regard this man as unwise in following the preferences he has, but we cannot regard him as mistaken in having them. For there is no such thing as being mistaken <u>in</u> one's preferences (although, as mentioned in Chapter One, a man may be mistaken about what his preferences really are). But if there is no such thing as being mistaken in one's preferences, equally there is no such thing as being <u>right</u> in one's preferences. One just <u>has</u> them, and that is that. Because of this, reason-giving, in the sense of justificatory explanation as opposed to clarificatory characterisation, has no role to play in our discourse about our tastes or likings or preferences.

However, reasons have a quite different and quite essential role in evaluation strictly so called. By 'evaluation' I mean any process of judging which makes implicit or explicit reference to standards which others can be expected to recognise and may well share, and which I intend to argue they normally <u>must</u> share. The judgement that a man is ill-mannered, or that he is strong, is an evaluative judgement in the sense I mean; for although these judgements signify a standard that is obviously relative in various ways to the context in which they are made, what that context is, and what standards consequently apply, will determine to a more or less exact degree what does and what does not count as strength

or bad manners within that context. Someone who disputes a judgement which is determined in this way can only escape the charge of being wrong in his judgement if he can show that the standards imposed by the context are in some way inappropriate to the sort of judgement he is making. For instance, a certain action might count as rude in the context of a formal dinner-party according to the standards which people generally accept and which are no doubt codified in various works on etiquette. If that is the context, and those are the appropriate standards, then a defender of that action must show that it served to fulfil the aim or function of having rules of etiquette — dictating forms of behaviour, so that everyone is put at ease — better than a slavish adherance to the rules would have done in those circumstances. Only in that way can he hope to show that the action was not, after all, ill-mannered. Otherwise he is wrong, and it was. Moves of this sort occur frequently in moral arguments too, where the question is about what someone ought to do or whether someone was right to do what he did. Only judgements about prima facie moral goods are immune from such considerations, just because they are context-free and in that sense absolute, since they are relevant whatever the context of the discussion.

The distinction between evaluations and expressions of preference is a sharp one, although its sharpness is obscured by the fact that there are many personal preferences that are widely held, so that it is possible to express a preference using a term which is suggestive of there being a standard, whether or not there actually is one. This seems to be the principal function of the word 'nice', for instance, a term which hints at there being some reason for approval, but from which it is particularly easy to retreat into the position of saying that one only meant that one liked the thing in question. One way in which the sharpness of the distinction is manifest is that there is a requirement for evaluations to be consistent with one another, whereas no-one is logically required to be consistent in his tastes. The requirement of consistency in evaluation is, of course, inherent in the

notion of there being standards involved, since nothing which was inconsistent could be said to adhere to any standard. But the other, and for our purposes more significant, basis of the distinction between preferences and evaluations is that the latter, but not the former, are always open to the significant request for a reason to be given as to why one has judged as one has.

It is always legitimate to ask why someone has appraised something as he has, when his appraisal amounts to an evaluation, because a genuine evaluation must always be based on reasons, whether or not the person making the evaluation is able to give them clear expression. The requirement that there be a reason rests on the fact that evaluation functions in a particular way, a way which Mrs Foot's analysis has done a good deal to make clear. For the function of evaluation is to make it clear how, or whether, something relates to our aims, purposes or aspirations; and as we have seen, the presupposition that we agree in our aims and also in what counts as achieving or furthering those aims makes it possible for the facts to determine our evaluations. And my argument is that those aims <u>have</u> to be shared if the whole process of demanding reasons and giving them is to have any point. If our aims, purposes and aspirations are <u>not</u> shared, then the whole edifice of moral <u>reasoning</u> collapses, and almost all of our moral discourse is revealed as a sham. But it is not a sham. Since in this respect moral discourse is just like any other type of evaluative discourse, to deny that there are good reasons for moral evaluations entails the denial that there are ever good reasons for evaluating anything at all.

Suppose I am discussing cricket with somebody, and I say that Smith is a good bowler. If I am asked why I say this, it is not an adequate response just to say that I like Smith's bowling. My use of the word 'good' has committed me to a form of discourse in which the giving of reasons is obligatory. Now suppose my reply is that Smith is fast and accurate as a bowler, and this is my reason for calling his bowling good. Now this <u>is</u> a

reason, because these facts about Smith satisfy criteria for being a good bowler; and my answer will also satisfy the questioner providing he knows enough about cricket to realise the significance of being fast and accurate. But suppose he does not realise it. He may demand a further reason for my appraisal of Smith as a bowler. The demand for a further reason can be met by explaining the point of a game of cricket to him, making him see that the faster and more accurate a bowler is, the more likely he is to succeed in what he is trying to do, viz. contain and dismiss the opposing batsman. If my questioner asks why Smith or any other bowler should be trying to do that, there is no further level of explanation available to me. For that just is what one tries to do if one is a bowler in cricket. The problem for the questioner, if he has understood all I have said, is not that he does not understand cricket; for I have explained cricket to him. What he does not understand is that when one has given the basis for an evaluation by relating what is being evaluated to the point of the event, state of affairs or activity of which the thing or activity being evaluated forms a part, one <u>has</u> explained one's evaluation. Two points of significance emerge from this. One is that evaluative judgements are impossible outside a context which includes a recognition of the thing or activity evaluated as having a certain point or aim. The other is that the explanation comes to an end when the point or aim has been stated and its relationship to whatever is being evaluated has been made clear. The point or aim cannot itself be evaluated, unless it is seen as a means to something further. And there may not be anything further. A good bowler would still be a good bowler if, as is sometimes suggested, cricket is itself an activity entirely lacking in any point or value beyond those implicit in the aim of the game. If the point of cricket were to change (say, by ceasing to be a game one tries to win and becoming instead an aesthetic ritual which would be spoiled if a wicket were ever to fall) then of course the criteria of good bowling would change too in response to the change in the point of cricket. But without there being a point of some sort to the activity, there could

be no such thing as an evaluation of anyone as a good or bad cricketer. For unless the game has a point, there cannot be criteria for evaluation, or reasons for evaluation given in terms of these criteria and their satisfaction. And reasons are what distinguish evaluations from expressions of liking or preference.

The same process of reason-giving occurs if, instead of saying that Smith is a good bowler, I say that he is a good father. Again, if I am asked why, it will not do to reply that I just like Smith's way of carrying on as a father. Having called him a <u>good</u> one, I must justify my ascription. I am quite likely to do this by pointing to certain facts about his behaviour and attitudes towards his children. These will, after all, be the kind of things I have noticed about Smith if I formed my opinion of him as a father for myself; or they are the kind of things I will have been told if I have my opinion from others. I might point out the amount of time he spends with his children, characteristic acts of generosity towards them, the way he is always willing to help with their homework, the sacrifices of time, money and patience he makes in their interest, and so on. If the questioner understands enough about being a father to appreciate why I mentioned these things, rather than mentioning beating, starving or totally ignoring the children, he will appreciate that the things I have mentioned satisfy criteria of goodness in fathers, and that will be sufficient to explain why I called Smith a good father. But if for any reason my questioner is dissatisfied with my answer, and demands a further reason for calling Smith a good father, I must explain to him what the social role of a father is, viz. to do all he can to care for his children so that they will develop into flourishing adults. And I will then say that the facts I have mentioned about Smith's conduct towards his children do in fact count as significant instances of such caring and are likely to have the anticipated result. If I am now asked <u>why</u> a father should care for his children and want them to grow into flourishing adults, I am hard put to find an answer, because I have already referred to the point or aim of fatherhood as a social role, and that was

the basis of my evaluation of Smith. I can only reply that that is what is expected of fathers, and Smith is, after all, a father and it was as such that he was being judged. For given that a person assumes or acquires the role of a father (whether by choice or not, whether willingly or unwillingly) his caring for his children, or his failing to do so, just is the respect in which he is a good or a bad one. Furthermore, since we know what the aim or point of being a father is, the question of what most effectively achieves that aim or point may be a matter open to dispute, but it is also in principle a matter of fact. So, unless I am mistaken about that, only Smith's doing or being the things I have mentioned, or things like them in that they will have the same result, can count as reasons for calling him a good father.

(I am not suggesting that the question 'why should a father care what happens to his children?' is necessarily a senseless one. It might be a not particularly clearly expressed query about the role that fathers are taken to have. Why should a father be more concerned for the welfare of his own children than the welfare of the children of others, or the welfare of others simpliciter? This may be seen as a question about the relative importance of different moral aims, in which case it is certainly a moral question, but falls beyond the scope of the present analysis for reasons given at the beginning of this chapter. Otherwise, it may be seen as a valid question about the social psychology or social history of the role of fathers. But then it is not a question of moral justification at all, but a question about the nature or origin of a social institution which provides part of the framework within which morality must function. Morality may lead us to question or challenge part of that framework; but later arguments will make clear that we can only do so from inside that framework.

But if the question 'why should a father care for his children?' is a request for a further level of enlightenment about what might be called the well-springs of morality — about whatever it is that makes people have any non-prudential attitudes at all — then there are two

possible responses we can give, according to how we understand what the questioner is puzzled about; and neither answer is likely to be the one the questioner wants. If the question is a first-order one, about what, in the final analysis, provides the impulse to moral action, then it is a request — and perhaps rather a pathetic one — for an insight into those feelings of human concern for others which the questioner somehow lacks or fails to understand. It is rather (though the analogy must not be taken too literally) like the blind man's questions about visual experience, the request for understanding about a faculty that can only be properly understood if one does not lack it. If, on the other hand, the question is a second-order one about the philosophical analysis or explanation of the concept of goodness over and above what has already been given, then the answer to be given can only be a descriptive one in which we indicate with a greater or lesser degree of sophistication a whole range of complex facts about human nature as it appears to us, and hence about the contexts in which humans do speak of moral goodness. It will not reveal any final justification of moral goodness _per se_, which is what the questioner presumably wanted. To frame an answer just by describing how we are constituted is not to lapse into a species of Intuitionism or to refer to what all men want, or even to what all men need. It is not even exactly to recognise what Rashdall called "tendencies to particular emotions, spontaneous tendencies to approve of certain kinds of conduct and to disapprove of others, which rest upon no logical ground, but must simply be taken as data upon which the Practical Reason has to work."* For the data Rashdall refers to is not exactly _raw_ data either, but data to be understood only in terms of the place which our whole complex pattern of moral thinking and feeling gives it. The descriptive answer is appropriate because we are at the point where explanation has come to an end. It is a requirement of there being moral judgements that people should have _some_ aims, should see the point of doing _some_ things, and

* H.Rashdall 'The Theory of Good and Evil' (London 1924) vol.1, p.161

have some feelings or desires which are non-prudential. But why we have what we have is not something that moral theory could be expected to justify, but only to describe. A fortiori, we cannot answer the questioner's final demand by giving a moral reason; for, as I have argued, moral reasons presuppose the existence of these aims, points, feelings or desires, and cannot therefore be expected to explain them.)

The examples of evaluation relating to fathers and bowlers which I have discussed are valuable because they indicate clearly the role of the person and hence the respect in which he is being evaluated, which makes it easy for us to see exactly what is to be expected of a person who has that role. This in turn enables us to see how evidence and reasons for evaluation can come into play in response to that expectation and how, in terms of that expectation, they can be regarded as determinate. Most ascriptions of moral value, however, are not, or are not primarily, dependent on a man's having a particular role, but depend just upon his being a man with the normal capacities to think, feel and act. Yet I maintain that this makes no essential difference to the structure of the evaluative processes we employ when we evaluate what he does or what he is, but merely makes it less easy for us to see clearly the connection that exists between the non-prudential aims we expect him to have as a man and the evidence we find relevant or decisive for our evaluations of him. Nevertheless, there are some very general aims or clusters of aims which we do not find it at all hard to state as conditions for moral goodness. We expect a man qua man to be concerned to relieve the suffering of others, to respect life and property, to be merciful, honest, considerate and just. Not all of these expectations are on the same level of generality, some being subsumable under others, not all are compatible in every case, and all are sufficiently general to allow a good measure of latitude, both within societies and between societies, as to what exactly is to count as an action which fulfils or tends to fulfil one of these aims. Nor do I claim this list is complete or incontrovertible. The openness of these expectations, which has sometimes been offered as a criticism

of descriptivist ethics, is no impediment to my argument, since, unlike Mrs Foot, I am not trying to claim that there are any moral aims which a man must have, or that there need be any particular content common to any action which could count as moral. I am only arguing for the logical point that, whatever these aims are and whatever counts as fulfilling them, there must be agreement as to what it is if the processes of evaluation are to occur.

What is called a moral principle is in fact no more than an embodiment of such an aim. The action-guiding function that moral principles have in our discourse requires that they be couched in prescriptive language; but that should not be allowed to obscure the fact that a moral principle could only be one if it expresses some non-prudential aim that a man might reasonably be expected to have, so that the state of affairs that the principle enjoins us to bring about if we can is one that we can see the point of, as having some impact on the welfare of others or some bearing on our general concern for others. Thus 'one ought to shout loudly on Tuesdays' could only count as a moral principle if shouting on Tuesdays was something we saw as worthwhile in itself, and moreover worthwhile in a way that bore some relation to the welfare of others. (The principle suggested does not strike us as a likely candidate to be a moral principle, no doubt; but I do not think we can rule it out a priori. For moral principles are contingent on our having the non-prudential aims we have, and so also is our having the structure of moral reasons we have.) Now it follows from what I have said earlier that a moral principle is not itself something that we can give further reasons for holding, since moral reasons are consequent upon our having certain aims, and hence these principles. The only justification that can be offered for a moral principle is the strictly logical one that it is deducible from another moral principle of greater generality which we hold — also without justification being available. There is one other respect which, it has sometimes been suggested, can lead to our talking about the justification of a moral principle, and that is by reference to our disinterested desire to see certain states of affairs brought

about. Thus it is sometimes said that a principle such as 'stealing is wrong' can be justified by the undesirability of the state of affairs in which a man is deprived of his goods. This is a harmless and possibly an enlightening line of argument, but it should not be mistaken for a moral justification. It is no more than the logical consequence of moral principles being expressions of our non-prudential aims. For an aim must logically involve a goal, which will be some state of affairs, so a principle must logically do likewise. To say, then, that a moral principle is justified by the states of affairs it would bring about is to say no more than if we said that hitting a target is a way of justifying having aimed at it. The only advantage of referring to desirable or undesirable states of affairs rather than moral principles which enjoin or prohibit is that principles wear an air of indefeasibility, whereas to recognise that certain states of affairs are good or desirable in themselves is no bar to our seeing that such states of affairs may be incompatible with other states of affairs which we also recognise as desirable. And it is my claim that the moral processes I am describing lead to <u>prima facie</u> moral principles, not to absolute ones.

Now it is sometimes suggested that anything at all may be proposed as a genuine reason for a moral valuation. For, it is said, a man can choose what he will value and why, and to deny this is to deny his essential freedom as a moral agent. The moral process is seen as one of deciding what has value and why, and commending it to others as well as letting it guide one's own actions. But this conception of moral freedom involves a confusion between the logically possible and the rationally coherent. It will be clear from my argument so far that anything at all could, logically, come to count as a reason for valuing something, just as anything at all could, logically, come to count as an aim that people have and consequently count as something they see as having intrinsic value. But while there might be no reason in that sense why people should not choose to value what they wish, for whatever reasons they wish, and see whatever point they have a mind to see in things, we cannot expect any of this to be intelligible to someone

who does not happen to share that particular view of things. By saying that a person may call anything he likes a reason for valuing something, we undermine the very notion of giving a reason. Consider the case of a man who claims to value a knife as an instrument for cutting, and when asked why he thinks it is a good one replies that whenever he uses it he is reminded of his grandmother. The words he utters are not unintelligible in the sense of involving a misuse of words, or generating a contradiction, or of not conveying a meaning. The words are unintelligible (as they stand) insofar as they are supposed to constitute a reason for his valuation. For a reason can only be a reason if it has some bearing on the case, and can only be a good reason insofar as it provides grounds for the judgement. And a reason can only be said to have a bearing on the case, or to constitute grounds for judgement, if it does so for other people, and not just for the person who utters it.

It is possible that a bizarre 'reason', like the one offered in the example just given, might be made intelligible by some further form of explanation. To take an example around which discussion has revolved, suppose a man claims that he has a duty to clasp and unclasp his hands.* We should find that puzzling until someone explains that the man is a sculptor who fends off rheumatism in order to finish works of art which he believes it to be his duty to complete for posterity. We can see that a conversation of the form:

(A) He feels a great sense of duty to posterity.
(B) How do you know?
(A) Well, just look at the way he clasps and unclasps his hands.

is distinctly unconvincing until the above explanation is added. But while this example shows that there are circumstances in which a fact as odd as a man's clasping and unclasping his hands could be a reason for commending him, it does not show that anything at all could count as a reason for commending or for valuing, and still less that it is open to us to commend or value anything we want to while remaining in-

* see R.M.Hare 'Descriptivism' (PBA 1963) p.129 ff.

telligible to others in the process. We are only able to appreciate that clasping and unclasping one's hands might be a morally good thing to do because we value art and recognise that on that account an artist has a duty to produce it, or at least is to be commended for doing so. We might well understand why an aging safebreaker engaged in the same exercises in order to keep his fingers supple for his trade, but we would find the suggestion that he had a moral duty to do so unintelligible, unless we happen to believe that theft is a morally good thing to engage in. There is all the difference in the world between seeing why somebody does something and seeing what he does as a reason for commending him. And the difference is that only where a moral principle is held, and a fact is accepted as a circumstance contributing to the fulfilment of the aim embodied in that principle, can that fact count as a moral reason. This, of course, is a direct consequence of the relation between aims and the facts which constitute or satisfy criteria for their achievement, which I discussed earlier in this chapter. I am not suggesting that Hare is unaware of any of this. His point in the paper from which the hand-clasping case is taken is only that the relation between moral beliefs and the facts which can be cited in support of them is something for us to determine, and not something imposed by the logic of unavoidable definitions and morally neutral facts. With this I agree. But I think Hare has failed to appreciate the significance of the fact that reason-giving only becomes possible once we have decided what our moral aims are, and what facts we take as evidence for their fulfilment. Once that has been done, we can give and understand reasons as reasons within a community in which everyone else has taken the same decisions about these matters as we have. But the consequence of that is that our moral aims, and a general sense of what is to count as achieving them, must be shared by all of us who are able to appreciate these reasons as reasons. Unless this is so, anything that is put foward as a reason for valuing something will fail to convince, not because it is a poor or inadequate reason, but because it will not count as a reason at all.

The necessity for there being agreement in the principles which embody our moral aims can be easily seen if we consider what would happen if everyone chose his own principles and no-one chose the same ones as anybody else. If that were to happen, there would of course be no agreement on criteria either, because of the way these are related to the aims one has and hence to the principles that express those aims. Different people might regard the same facts as criteria, but they would not be criteria for the same principle, so they would not be the same criteria. But in that case, no person could ever give as a reason for a judgement anything that would count as a reason to anyone else; for to give a reason is essentially to show the relationship between a fact and a principle via a criterion. and ex hypothesi this could never happen. But if for no-one except the giver of it is a reason ever a reason, the notion of reason-giving could play no part in moral discourse — but then of course there would be no such thing as moral discourse anyway, since what characteristically distinguishes evaluation of any sort from expressions of preference and likings is precisely the fact that reasons can be, and need to be, given.

It might be thought that if my account of our moral thinking and discourse is correct, there is very little room left for moral disagreement. This is not so, however, and for more than one reason. The main reason, which concerns the possibility of moral subjectivity, will appear in its proper place in a later chapter. However, even here we can see that there is plenty of room for moral dispute since there is plenty of room for dispute about what the facts are and what precisely is to count as a criterion in particular cases. And here, it will be remembered, we are still only talking about the justification of prima facie principles, and not about conflicts between principles at all. As an illustration of this, let us turn to an example I mentioned earlier of the father who sacrifices his children to the gods, not as an act of piety but as an act of caring. I only said then that this would involve the father in some rather unusual beliefs. We can see now that for his act to be a caring act it would have to be true that his

children actually benefited from it, and were assuredly better off 'living with the gods' than they would have been living on Earth with their father. If the father believed this wholeheartedly, we can imagine him sacrificing his children, sadly no doubt, but believing that he was acting as a good father should. It is unlikely that we would act as he did, but that is because we do not share his beliefs, not because we do not share his principles or lack the courage of our convictions. For in fact his principles are like ours in that he wants only the best for his children. And it is interesting to note that, if his belief were true, he would have hit upon a uniquely decisive criterion for being a good father since it would guarantee to fulfil our expectations of what a father should aim for in bringing up his children. This would not discredit all the more usual criteria of being a good father, but it would place them in a new perspective. This macabre and not over-subtle example shows just how much scope there might be for differences of practice and a different appreciation of criteria even though we share the same principles.

It might also be objected that this account of reason-giving in moral discourse is mistaken because the notion of a moral reason is itself suspect. This amounts to the idea that what we call reasons in moral discourse are not reasons at all, but only resemble them in giving the appearance of justification. In that case, since we evidently do go through processes called 'giving reasons for moral judgements' and 'being convinced by moral reasons', there must be some practical connection between the giving of reasons and the holding of moral convictions. Let us consider the Emotivists' suggestion that most sorts of moral reasoning amount to persuasive mechanisms with a causal but not a logical capacity to produce moral beliefs (or, as Emotivists put it, to change or foster attitudes). On this account, then, a moral reason would be an effective mechanism for altering, fostering or bolstering our moral attitudes, but this mechanism would only be effective insofar as it operates on those who are predisposed to be affected by it. And this implies that the people for whom

the proposed line of argument constitutes a reason must already agree with the proponent of the reason in their outlook. So much Emotivists would claim; and so far their position is not at variance with my own contention that reason-giving demonstrates prior agreement in moral principles. The sting to the objection is in its tail; for it is then objected that this basis for a reason being able to count as a reason is to be contrasted with reason-giving backed by the force of logic. Presumably what is meant by the notion of a 'real' reason here is that it is one which no rational man could fail to accept, by virtue of reason alone. But surely this is an empty contrast! Moral reasons work as reasons because they make reference to moral principles; and these principles express the (non-prudential) aim a man has or point a man can be expected to see in some aspect of human activity or the human state. Without this having of some aim or seeing of some point in the first place, there is nothing for any 'real' reason to be a reason for. So there cannot, I suggest, be an argument which would lead a man, by the exercise of his rational powers alone, from the facts to a position of moral belief unless he already possessed some values which make it possible for the reasons contained in the argument to function as reasons — that is, unless he had already what Emotivists would call a psychological predisposition to accept that certain facts will count as reasons for him. The activity of pure rationality in deriving a moral conclusion that moves one to act out of a body of facts alone without other assistance is, of course, pure fantasy. Once we realise that this is what is being contrasted with what Emotivists call the psychological propensity to accept something as a reason, we see that the only threat it poses to my theory is that of calling the justification of the reason a causal or psychological one. For this is misleading, inasmuch as (a) it makes moral reasons seem insecure as justifications rather than as persuasive mechanisms, and (b) it allows that what might count as good grounds for a moral belief in one culture or society might not in another. (I am not here suggesting that there is no such thing as a distinction between good and bad reasons in moral discourse. There is.

But I do deny that a moral reason cannot be a good one unless it can be shown to be free of all influence from moral beliefs and universally applicable by virtue of reason alone.) The Emotivists were not mistaken in thinking that value-free rational thought alone cannot justify value judgements. However, they were mistaken if they believed that only in that way could moral judgements be capable of interpersonal justification sufficiently strong for us to be able to deny that they are merely subjective.

My argument, then, amounts to saying that reasons will count as reasons only given the existence of a moral standpoint. And they are only communicable as reasons if that standpoint is a shared one. Thus I am saying that whatever can count as a moral reason must be seen as such from within a system of shared moral beliefs. And when it <u>is</u> seen in this way, I claim, it will fulfil the requirements for being a criterion and so render the moral judgements which fall within its scope determinate. This creates conditions in which, in the everyday sense of the word, it would not be unreasonable to call moral judgements about what is good objective; although on the analysis given so far, such judgements do not yet meet the strict conditions laid down for objectivity in earlier chapters. For we have not yet established whether there is just <u>one</u> moral standpoint or whether there are <u>many</u>. Nevertheless, the rightness of a judgement according to the account we have given so far will not depend upon the attitude of the individual judging, but upon the construction which, if he is reasonable, he <u>must</u> put upon the facts according to the moral outlook from which he judges. And it is this outlook, as expressed in the shared moral principles he holds, which gives the facts relevance and makes the giving of reasons intelligible as nothing else could.

As a final point, I must emphasise that this notion of a shared moral outlook must not be construed too narrowly. To say that one stands within it does not mean that one will necessarily agree with all or most of the judgements that others make about what is right or what is obligatory. It does not even mean that one will <u>normally</u> do so. For that will only occur if those around one also share the same <u>moral code</u> as oneself. One's moral code is determined

partly by the conditions imposed by the environment in which one must live, and partly by the social institutions which have arisen in response to that environment and man's need to live with others in it. It is at this level of moral code that morality can be described as a range of experiments, or hypothetical solutions to the problems of communal life.* But to say that morality is only rational given a shared moral standpoint is not to single out any particular moral code and set it on a pedestal. It is, rather, to point to a range of human concerns which every moral code must be seen as an attempt to face up to and express; and so, in a sense, it is to give a meaning to the term 'moral' by indicating what that term covers. No doubt no moral code will succeed in covering that area of concern completely, or without its own inbred bias. But all such codes will, to a greater or lesser extent, share the same general aims. And insofar as people are prepared to be rational about their beliefs, we may take this common stock of aims as the basis for meaningful discussion, even across moral codes, of how a man ought to live. My having referred to a moral standpoint, then, is in a way misleading if it suggests that there could be more than one of these, the standpoint to which different moral codes will give the fullest expression they can, incomplete though that is always likely to be. To say that there might be more than one moral standpoint is to admit that in the last resort a moral code is always a consensus, and that moral justification is always in the final analysis arbitrary. My object in the next chapter is to show that this cannot be so.

* for a view of this sort, see M.& A.Edel 'Anthropology and Ethics' (Case W.R.U. Press 1968)

CHAPTER EIGHT

EVALUATION AND MEANING

As I remarked at the beginning of Chapter Seven, the fact that we are able to support our moral conclusions with reasons only because we agree in the fundamental moral principles we hold is not in itself enough to prove that our moral principles are objective in the strict sense of the word. It only shows that moral rationality requires there to be interpersonal agreement in principles. It is still possible to say that a shared moral outlook, of the sort I have described as making moral reasons possible and intelligible, is only one moral standpoint among many other possible ones; and that insofar as we can call it a justified moral standpoint, its justification will only be justification for those whose standpoint it is. From inside it will appear to be objective; from outside it will be seen as nothing more than a consensus of moral opinions. In order to show that it is appropriate to call our moral standpoint an objective one, we will have to consider a further line of argument which will involve drawing together several ideas established earlier on in this thesis.

I argued in Chapter Two that the condition for there being sufficient justification for a method of judgement, and hence for the objectivity of the judgements made in accordance with that method, was whether the method generated unique and universally accepted judgements, so that the method provided the only correct way to judge. The test, therefore, of whether a moral standpoint is only one among many, or whether it is the only correct basis from which to judge, is the test of whether it is possible to step outside it, and see matters in some other, and equally satisfactory, way. If it is possible to do this, then moral objectivity does not exist, though we can understand why we might think it does, because our moral dealings with others who share our standpoint would create that impression. If, on the other hand, it is not after all possible to step outside our moral standpoint, then the fundamental principles which together constitute that standpoint must be objective.

However, it will have undoubtedly occurred to the reader that this capacity to step outside any moral standpoint, and reflect on it critically, is precisely what has so often been claimed for the individual as the surest sign of his own moral autonomy and that of the human moral faculty in general. Thus it might be thought that in the previous chapter I have overlooked the existence of that capacity by describing our moral aims as if there were no distinction to be drawn between the good that men seek and what it is to be good (a charge that has often been levelled at Aristotle). For isn't it perfectly possible to step back from the pursuit of our moral aims and ask whether what we call good really is good? Indeed, isn't it vital for the health of morality that this process of reflective detachment should occur? (I call it detachment only because it involves separating oneself from one's conventional beliefs, and not because it represents no involvement in the world. To the Existentialist, this agonising process of detachment seems to be the only authentic way of coming to grips with the world.) It is not my intention to deny that we can, and sometimes should, stand back from our beliefs. I only deny that we can stand back from all of them and so move outside the moral standpoint altogether in order to create, or recreate, our own. To think that we can do that is to confuse the capacity to decide what is good with the imagined capacity to decide what goodness is to be in one's own case.

Within the tradition of British moral philosophy, the most recent important example of a theory which enshrines the individual's freedom to call all existing moral beliefs into question is to be found in the work of Prof. R.M.Hare. In the terms of his prescriptivist theory, it might be said that in my outline of moral justification in the last chapter, I have failed to recognise the important distinction he draws between the descriptive meaning and the prescriptive meaning of moral terms, and so ignored a crucial aspect of the moral situation, the individual's (rational) capacity to decide on his own moral standards. According to this view, it is possible for someone to understand the whole structure of the relations between the principles we hold and the sort of facts which we take to be morally significant in the light of those principles, and yet to refuse

point blank to call any of those principles his own. In other words (and taking 'good' to be the term characteristically employed to express prima facie value) he might accept that we give a certain complex set of descriptive meanings to the word 'good' in various moral contexts, and commend whatever fits any of those descriptions; but although he can understand exactly what we are doing (though he will not understand why), he is not prepared to commend anything we commend, because the descriptive meaning he thinks 'good' has is entirely different, and so 'good' will be used by him in entirely different circumstances. According to Hare, primarily evaluative terms such as 'good' have no particular descriptive meaning attached to them. Their meaning resides in their being used to commend according to whatever descriptive meaning they happen to have been assigned by the person using them. So there can be no reason for saying that what we think morally good has any better claim to be really good than what somebody else with completely different principles thinks good. For the essential meaning of 'good' is the same in both cases.

In fact it is not quite true to say that in the last chapter I entirely ignored the possibility that a man might not share the moral standpoint of those around him. For I did refer at one point to an occurrence that covered this case. I suggested that someone might want a further explanation of a fundamental moral principle because he was not convinced by the previous explanations — those which laid stress on facts that satisfy criteria for moral worth, and so should lead him to a position where he can recognise the principle for what it is, an expression of an aim he is to be expected to share. This allows for the possibility that his failure to acknowledge the principle as a principle for him is owing to his having a wholly different moral standpoint. I described his failure there as a lack of understanding, because that's what it usually is amongst those who share our outlook in other respects; but I was also to some extent anticipating the arguments of the present chapter by implying that if his standpoint is a wholly different one from ours, he will still lack understanding, but in a more radical fashion. We can see this easily enough. A man who fails to see the significance of

some straightfoward fact and be moved by it — that doing such-and-such will relieve the suffering of another, for example — so that he fails to appreciate that doing it should be an aim for him, has a less than complete understanding of the concept of moral goodness. A man who is never moved by the significance of any fact of the sort that moves us in this way lacks any understanding of that concept. It will not do to say that he has a different concept of moral goodness. For whatever he has (if indeed he has a concept at all here) will not be anything we could understand as moral goodness, so it would be a mistake to use the expression 'moral goodness' to describe it. That all this is so will, I hope, be made clear in the course of this chapter.

Prescriptivism represents an attempt to square the moral freedom of the individual with the notion that moral beliefs can be held on rational grounds; and one of its most striking features is its emphasis on the ability of the individual to step outside any existing moral standpoint and establish his own personal one. Indeed, if prescriptivism is strictly interpreted, 'the moral standpoint' could only ever be a consensus view which must be regarded as an arbitrary consensus when seen from the outside. Its claim to be anything more than this rests on Hare's insistence that to be a moral standpoint a person's principles must meet a series of formal requirements. Principles must be prescriptive, universalised, and maintained with logical consistency by their holder, so that he will not balk at the idea that the prescriptions he makes for others must apply equally to himself. Prescriptivism still allows for a plurality of arbitrary standpoints, though, in spite of the fact that much of Hare's energy has gone into trying to prove otherwise, especially in his second book, 'Freedom and Reason'. His argument here is essentially that what people think others should or should not do will be subject to continual checks and modifications as they reconcile their universalised beliefs with their own inclinations, when put in real or imaginary situations in which those beliefs apply to themselves. Only fanatics, says Hare,

will be proof against this process of moral modification and accommodation. I think that Hare overlooks the real significance of the fact that fanaticism is moral courage viewed from a different standpoint; or that at least, if morality is pluralistic in the way his theory allows it to be, we cannot deny that the fanatic's position has as much right to be called a moral one as our own, and that our own position is as open to the charge of arbitrariness as his is. In the end, Hare's suggestion that our moral standpoint as represented by his theory is not an arbitrary consensus boils down to a faith in one particular set of inclinations that people might have rather than in others. Yet many attitudes that a person might adopt would satisfy Hare's conditions for counting as <u>moral</u> attitudes, although we would in all probability not want to regard them as views which could be morally justified. For example, the man who consistently exacts a savage penalty on others for what we would regard as an innocuous act might indeed not want that savage penalty to be applied to him. But he might have no aversion, in the light of his principle, to its being exacted on him-as-a-person-who-could-do-such-a-thing. It is not that he logically cannot imagine himself doing such a thing, but simply that he would regard it as grossly immoral to do it. And on Hare's account of morality, we cannot accuse him of any failure of moral understanding.*

Prescriptivism has frequently been attacked in this way for having the consequences it has. However, I intend to criticise the theory from the other end, so to speak, by attacking its initial presuppositions about the nature of moral meaning. For the fact that prescriptivism allows the individual such complete freedom to adopt his own moral standpoint is owing to the doctrine that the complete separation of prescriptive and descriptive elements within the meaning of general moral terms is compatible with their having a meaning at all. Hare's adoption of this view is the result of his misinterpretation of the significance of the (logical) openness of moral arguments which I discussed in Chapter Six. There I argued that the fact that 'X is good'

* for more detailed discussion of the failure of this part of Hare's theory, see T.D.Perry 'Moral Reasoning and Truth' (Oxford 1976) pp.173-84, and R.B.Brandt's review of 'Freedom and Reason' in the Journal of Philosophy (1964)

is never true by virtue of the definition of X, shows no more than that 'X is good' is contingent. Hare takes this contingency to be a sign that the individual is free to form his own conception of the evidence and reasons that will count in favour of the goodness of X; whereas I maintain that the contingency shows that the question 'is X good?' must be decided on what the evidence shows according to some public standard of what is to count as evidence. As I suggested in that chapter, it is an inconsistency to treat empirical judgements in one way and evaluative judgements in another just because evaluations are supervenient on other judgements, and so involve a further stage in reasoning, which we often cannot perform conclusively, and so come to the impression that moral judging is a more open process than it really is. In Chapter Seven I argued that to treat evidence and reasons as self-generated by the individual makes the notion of a moral reason incoherent. I will now argue that it also makes incoherent the notion that a primary moral term (as Hare describes it) could have a meaning. If this is correct, it follows that prescriptivism, and with it the belief that a person can stand wholly outside the moral standpoint to criticise it, is broken-backed. For in order to stand outside the existing moral outlook to criticise it, one must be able to give meaning to the terms in which one frames one's criticisms. And that, I claim, is impossible to do.

Hare distinguishes two elements which combine to make up the meaning of an evaluative term (although he sometimes speaks of the dominant element as being <u>the</u> meaning of the word.) What he calls secondarily evaluative terms have fixed descriptive connotations, and the properties connoted constitute criteria for the correct use of the term. There is also a prescriptive element within the meaning of these terms, which indicates our propensity to commend things for having those properties. But we are always free to drop the prescriptive element from the meaning if we no longer wish to commend anything for having those qualities. Thus (to take Hare's own example) to call a man industrious is to describe him as hard-working, but also to praise him for it. But because here the descriptive meaning is paramount

we should continue to call a hard-working man industrious even if our attitude to being hard working were to change so that we now used the term in a neutral or even pejorative way. The prescriptive meaning would have been lost or been changed. (In fact, it is far more likely that we would use a different word, such as 'workaholic', if we wanted to signal a change in our attitude. The attitudinal connotations of a word are not very easy to shed. But I do not think this affects the substance of the distinction Hare is drawing.)

Words which Hare calls primarily evaluative terms, on the other hand, have a dominant prescriptive element in their meaning. 'Good' is Hare's foremost example of such a word. Here it is the attitudinal element in the meaning that is constant, and it is the descriptive element — what factual qualities the word connotes — which can change. Now this change can take place in two quite separate ways, which Hare is perhaps not ready enough to keep separate. On the one hand, a term of great generality of application, such as 'good', changes its application according to what sort of thing is being commended by calling it good. There is, after all, no property that a good fire-extinguisher has in common with a good deed in virtue of which we call them both good. But it does not follow from this that there are no criteria for being a good fire-extinguisher or for being a good deed, but only that the appropriate set of criteria come into play according to what is being evaluated. The other way in which a primarily evaluative term can change the descriptive element in its meaning is that people can change their minds about what things are good, or in what respects things are good. That is, opinions about what properties in the world are good-making criteria for this or that can and do change. It is by changing the descriptive content of the meaning of 'good' while being able to keep its evaluative element constant, says Hare, that we are able to express changing standards of value. Both these pieces of analysis are valuable insights into the workings of evaluative language, and we are indebted to Hare for them. The problem arises from the use he makes of them.

Now Hare is sufficiently impressed by the extent to which words like 'good' can change their descriptive connotations in both these ways to come to regard such words as having meanings which are necessarily capable of being free of any particular set of descriptive connotations in any instance of their use to commend. This is more or less all right; and it is by no means confined to words which are used to evaluate. Very large numbers of words of all kinds systematically change their meaning according to the context they occur in — dictionaries list the different meanings and indicate the contexts; and any word which has descriptive connotations can have them changed if people come to use the word in a different way — the change can come about as a gradual shift in meaning or by stipulative definition in certain cases. But both these types of change require the usage of the word as between people to remain approximately constant, for that is the basis for saying that the word means something.

But because Hare is concerned with the meaning of 'good' and similar words, he is looking for a common feature of all uses of such words, and finds it in their use to express a choice. Just how apt the notion of a choice is for this task is not especially important. The important thing is that by seeking a common and constant feature of all uses of primarily evaluative terms, Hare has converted the ability of those words to free themselves from any particular set of descriptive connotations in any instance of their use, into the freedom to be separate from any particular set of descriptive connotations in every instance of their use. Thus every use of the word 'good', on Hare's account, involves a decision about the descriptive connotations it shall have in that instance, and hence on the criteria for its use. It is, I think, not always realised just how radical a proposal this is, for in dozens of examples that Hare mentions in his books, what he describes is our acknowledging an existing, publicly-accepted standard of merit for whatever is being talked about, or deciding between various relevant possible criteria on the same basis. To say here that one chooses to adopt these criteria, rather than any other logically possible ones, is

indeed to say that one chooses. But to choose in this way is to do something that can be done sensibly or foolishly, with or without a proper understanding of what one is doing. The instruction 'choose the shortest route' can obviously be followed correctly or incorrectly. It is less obvious that the instruction 'choose the best route' could be followed incorrectly, since it is always open to the person choosing to explain or justify any respect in which he decides to value the route he has chosen. But although there he may set up a new criterion by happening to think of something no-one else has thought of, he must get them to accept his new criterion as a criterion before it will count as one (or failing that, his criterion must be one that in general people of sufficient understanding would agree did count as one). And if he cannot do that, his proffered criterion does not contribute to a choice of the best route, and he cannot have chosen the best route on account of it.

Hare attempts to evade this difficulty, brought about by the need to make a criterion the subject of public agreement before it counts as one, by arguing that although the fact that a man has chosen X on account of C is not sufficient grounds for saying that X is good on account of C, it is sufficient grounds for the man's saying that he thinks X is good on account of C.* But this will not do. For where we can make the distinction, for any Y, between X's being Y and someone thinking or believing that X is Y, we require there to be a way of determining whether X is really Y or not. This, as we have seen, is a requirement for all empirical judgements, since there a distinction between belief and justified belief can always be drawn. And it is also true for any judgement where criteria which offer justification can have a place. Evaluations constitute such judgements, or we should never try to differentiate between them and expressions of subjective approval.

Yet it is really the expression of subjective approval that Hare is offering us here when he suggests that in calling something good, a person means that he would choose

* see 'Descriptivism' op.cit. p.127 and 'The Language of Morals' pp.107-8

it and choose it for the reasons he chooses to choose it. For while a person can make an evaluation on the basis of a novel criterion (if in general there are interpersonal criteria for this sort of choice, because the word 'good' has a measure of agreed meaning in this area of judgement, so that it could be said that a person had made a good or bad or sensible or silly choice here), that is not at all what Hare's theory is suggesting or allows. What he is suggesting is, as I said, something far more radical. For in making both the choice of what is valued and the choice of grounds for valuing it the province of the individual's judgement, as a consequence of locating the meaning of words such as 'good' within that double act of choosing, Hare has abolished the possibility of there being a distinction between deciding what is good and deciding what 'good' is to mean in the context of any particular act of evaluation.

This is a radical step indeed. Let us look at what it entails. We begin from Hare's insistence that in making an evaluative judgement it is always possible for the individual to step back from any pre-existing standard and ask himself whether he does in fact approve of that standard; and we must note that, according to Hare's conception of what it is to evaluate, it is not only always possible to do this but it is always necessary to do it. For on this view to evaluate is to indicate a choice, not only between possible objects of choice, but also between possible criteria for choosing such objects. If the individual is to be sure that he does really wish to commend something, he cannot merely adopt such existing criteria for judgement as he or his society happen to have to hand, but must either choose that these criteria shall be his standard here (in which case he literally re-creates a standard), or else must create a fresh standard of his own. Either way, it follows that every act of evaluation the prescriptivist performs involves a decision to adopt a standard he has chosen for himself. This is what is meant by saying that prescriptive meaning is always logically prior to descriptive meaning; any standard must (logically) be established before it can be invoked, and its establishment is (logically) an instance of the prescriptive use of the

value term involved. The standard the individual decides on must thus be implicit in his having chosen to value whatever he has chosen, even if he does not spell out what the standard is at the time. Although Hare often seems to be talking as we all do, as if criteria for evaluation were to be thought of as something independent of the individual's choice that he normally subscribed to in judging, according to his theory criteria must be called into existence afresh by each and every act of judging, and are therefore valid, so to speak, only for the duration of the act. For a fresh judgement will call for a fresh decision as to criteria. In this respect, prescriptivism is extraordinarily similar to existentialism, and one is reminded of Sartre's remark that through morality man is condemned at every instant to invent man.*

A second remarkable consequence of Hare's view of the meaning of moral terms is one that we have already noted in part. It is that the distinction, which we all draw, between 'P thinks X is good' and 'X is good' is simply not available for the prescriptivist to use. For, as we have seen, that distinction depends on there being an independent standard which judgements may meet or fail to meet; but on Hare's account every act of judging must dismantle all previously existing standards which could apply to the case in question. Other people may seem to think that this or that is good. But since they cannot have access to the standard of judgement involved — for that is something the individual creates by judging — and since the act-of-creating-a-standard-in-judging is what determines the meaning of 'good', we cannot actually tell what anybody else might mean by the word 'good', and so we cannot tell what they might be thinking.

The third consequence of Hare's view, which is a corollary of the other two, is that a man can never be wrong when he judges that something is good. For if, as I claim, 'X is good' is synonymous on Hare's account with 'X is what I would choose according to the criteria I determine in choosing it', it follows that every use of an evaluative term in judgement is as fully justified as it is

* in 'Existentialism and Humanism' (tr.P.Mairet, Methuen 1948) p.34

possible for the use of such a term to be. The reason is once more simply that the decision that something has value is also the decision as to what is to count as having value, so that in judging in accordance with the standard one sets by the act of judging (as one must) one must always judge aright.

It will, I hope, have become obvious from the foregoing that when the prescriptivist analyses the meaning of primarily evaluative terms (the ones the descriptivist's theory is supposed not to be able to cope with), he is able to make such terms wholly independent of any particular descriptive content and thus any pre-existing standard of judgement only at a very great cost. It is the cost of making words such as 'good' altogether incapable of having a meaning. For the double choice — of what to commend and what to commend it for — which a man is supposed to be making whenever he uses such a term, is, according to the prescriptivist analysis as I have presented it, a logically private occurrence. And the utterance of the word 'good' or some cognate of it would, on that analysis, be a speech act in a logically private language. The analysis in question therefore invites the full weight of the Wittgensteinian Private Language Argument against it. For as I have said elsewhere, that argument is not restricted to cases which involve inner sensations, but applies to any circumstance where the claim is that a person's understanding of a concept, and his consequent ability to communicate about it, derives entirely from his own case.

I shall not subject the reader to a second journey through the rugged terrain of the Private Language Argument, which was, I hope, adequately traversed in the first part of Chapter Five; for what was conclusive there is, mutatis mutandis conclusive here, and to give a logically private meaning to 'good' is no more possible than it is to give a logically private meaning to 'red', and for exactly the same reasons. I shall confine myself to re-emphasising just one point. It is not just true that a person who uses a word in accordance with 'rules' for its use that he alone devises cannot make himself understood to others. It is also true that he himself cannot claim to know what

he means when he thinks he is communicating with himself. What the misguided prescriptivist (and the existentialist too, so far as he shares that position) imagines himself to be describing in his references to the continually repeated act of judging *ab initio* which I have been criticising here is not moral judgement. It is not even the evincement of moral feeling that is being described. If it is anything, it is rather the raw material out of which moral feeling and moral understanding can be generated which is being referred to. Moral feeling and moral understanding can be, so to speak, elements within our experience only when they have been incorporated into a conceptual structure which makes moral thoughts and feelings intelligible *as* moral thoughts and feelings, and moral terms a means of expression of moral judgements within a public language. This, as Wittgenstein has shown, and as I have argued in his wake, is a process which the individual takes part in, but *not* a process which the individual can perform *for* himself or *by* himself.

Thus we can see that the individual is *not* able to step wholly outside the moral standpoint that he (and his society) occupies, and still retain an understanding of the concept of value which would make his alternative view intelligible. The moral critic must always criticise from within the system of beliefs he wants to alter or develop, no matter how revolutionary or incisive his criticism is. Criticism must always be piecemeal, never wholesale. To say this is not to say that a man may not justifiably feel that there is a very great deal wrong with the society he lives in, and that perhaps it would be better destroyed than reformed. This is a coherent position. To think it is not is to confuse the moral standpoint with a particular moral code which a society has and which, though it must lie within the moral standpoint, may be such a partial and distorted expression of it that reform is impracticable. It is as well to remember, too, that it is still *prima facie* goodness that we are concerned with, and that the differences which exist between societies are, for the most part, not differences about what is fundamentally

good, but about which of these goods are the most important.

Now the objection may well be made that in rejecting Hare's analysis of the meaning of moral terms, I am significantly misrepresenting his views. In one respect this must indeed be true. For most of the time Hare's views are at odds with themselves. Hare writes, for the most part, as if he is within the moral (and, generally, the evaluative) standpoint we all share, and which we are all able to dissent from or develop in various piecemeal ways. It is just that the theory of total moral freedom to judge which his analysis presents is wholly incompatible with that position. This need not surprise us. It is, after all, not very hard to discuss a word such as 'good' (or 'red') which we all know how to use, and to go on using it in the ordinary way while at the same time erecting a theory about its meaning which is incompatible with that use. Hare's books abound with interesting and rewarding examples which are only unsatisfactory if we also accept the full implications of the theory in whose favour they are set out. And Hare, I think, does not always appreciate what those implications are. For example, in 'The Language of Morals', Hare discusses a case where agreed public criteria for evaluation are unavailable because they do not yet exist, and notes quite correctly that we are able to use evaluative language in advance of there being such criteria. Now the fact that we are able to extend, change and develop the descriptive meaning of 'good' so readily is one of its most arresting features. What this does <u>not</u> mean, though, is that we can divorce the meaning of the word completely from <u>all</u> agreed criteria at once and still retain a grasp on its meaning. It is only if we do try to match Hare's example squarely against that suggestion that we come to see how absurd the suggestion is. Hare writes:

"Suppose that someone starts collecting cacti for the first time and puts one on his mantel-piece — the only cactus in the country. Suppose then that a friend sees it, and says 'I must have one of those'; so he sends for one from wherever they grow, and puts it on his mantel-piece, and when his friend comes in, he says 'I've got a better cactus than yours'. But how does he know how to apply the word in this way? He has

never learnt to apply 'good' to cacti; he does not even know any criteria for telling a good cactus from a bad one (for as yet there are none); but he has learnt to use the word 'good', and having learnt that, he can apply it to any class of objects that he requires to place in order of merit. He and his friend may dispute about the criteria of good cacti; they may attempt to set up rival criteria; but they could not even do this unless they were from the start under no difficulty in using the word 'good'." (pp. 96-7).

The eventual conclusion Hare draws from this is that knowing the meaning of 'good' "has nothing to do with criteria" (op. cit. p.110). Let us see if Hare's example really takes us (or him) in that direction.

When Hare says first that as yet there are no criteria for good cacti, but that it is still possible to talk of a good cactus and mean something, he is being somewhat disingenuous. For there are of course possible public criteria which will stretch to covering cacti as well, and which the speaker is likely to have in mind when he calls his cactus a good one. Because cacti are the sort of things they are, we would all have some idea about what the criteria for being good ones are going to be; at the very least, we know what they won't be — for instance, dead, smashed, shrivelled, obviously diseased. Knowing even this much, we can see that a tentative use of 'good' here is likely to be determined by these sorts of consideration. But even if we did not know this (perhaps because cacti were so unlike anything we had ever seen before) we would start casting about for parallels with things we all already know how to value. There are very few things we might conceivably think of valuing which have no parallels with anything we do not already know how to value according to the appropriate criteria. Whether the speaker has used the word 'good' rightly as well as intelligibly is something we will be better able to answer when the dispute about standards has got a bit further, when there are more people involved and more specimens available for comparison, and so on. The obvious point here is that the speaker is able to use the word 'good' meaningfully because there are already accepted

ways of extending more or less appropriate criteria to cases of this sort, and techniques for refining them subsequently. The speaker is not suggesting (and Hare does not imply that he *is* suggesting) that in speaking of his cactus as a good one he is also making a declaration of what the criteria for 'good cactus' *shall be*. He is only proposing what, given what things are normally taken to count as criteria in cases of this sort, the criteria for 'good cactus' *should be* (where 'should be' has something of the force of 'are going to be if I'm thinking along the right lines'). *A fortiori*, he is not declaring what the criteria for good cacti *shall be for him*. The link between the existence of public criteria and the meaning of 'good' in a given context would be most evident if the speaker were to break the link, by saying, for instance, 'mine is a much better cactus than yours because mine is dead and yours isn't'. If we took his remark seriously, we would be bound to wonder if he knew what 'better' means. We could understand his liking his own cactus more than another man's, perhaps, because it's his and he is attached to it. (This is the notion of sentimental value). But why does he say that his cactus is *better*? The appeal for a reason shows us immediately that interpersonal agreement in standards, and hence in judgements, is at some stage presupposed here. And we can see that Hare is by no means rejecting the mechanisms of interpersonal discussion and dispute here, because he at once suggests that the two friends could dispute about each other's proposed criteria for cacti, which means that at the very least they must be able to agree on what is relevant to their discussion.

Now suppose, however, that we take seriously the implications of Hare's contention that the meaning of 'good' and similar primarily evaluative words resides in a fixed use to commend plus a discriptive element which the person commending may determine for himself. Then — assuming for the moment it makes sense to say this, and remembering that 'may' here really means 'must' — the man who commends his cactus for being dead is right to think it good on account of this. (And even if 'may determine' did not mean 'must determine', we can see that he could still choose to think that way if he wished.) For here

'good' means his use to commend for whatever reasons (in accordance with whatever criteria) he chooses to commend. The man is not only using 'good' meaningfully according to the implications of Hare's theory of how evaluative terms mean things, but he is also bound to have judged correctly in praising his own cactus because he is determining the standard that he deems will apply at the same time, and in the same act, as he commends according to it. If the second man now calls his cactus a good one (say on account of its not having the feature which was supposed by the other man to make the other cactus a good one) then he too is right in his judgement, and for the same reasons. Whatever each man chooses to think right will be right, which only means that here we cannot talk about making right judgements at all, since the word 'right' has no function in the absence of anything that could count as being wrong. Any attempt to dispute about the criteria for good cacti is now pointless, as well as being impossible. It is pointless because the question of what a good cactus is like has already been settled for each disputant — at least, until the moment he comes to consider the question again. It is impossible because a dispute requires the possibility to exist that parties to it could be right or wrong. There must be something genuinely independent here for the dispute to be about. And here there is nothing. Everyone is right about his own understanding of the situation, meaning that here there is no 'right' and no understanding either.

From a comparison of these two interpretations of what is going on in Hare's example, we can see that what Hare really wanted from his example was evidence that 'good' could be brought to bear on new things without any change of meaning; and this is something the example is able to show. But it only shows it on the premiss (which Hare accepts in practice) that we all agree on what sort of things can meaningfully be called good in a variety of cases. And that amounts to saying that the meaningful use of 'good' to commend involves there being normal agreement about what things are in fact commendable and why, in established cases — something which will be clear from earlier parts of this

thesis, where I stressed the significance of the insight that if language is to be a means of communication, there must be agreement in judgements.* This (normal) agreement in judgements is quite compatible with the individual's freedom to decide, in any particular case, including new ones, what criteria should apply according to his (shared) understanding of what it means to call something good. Hare's example only degenerates into absurdity if we assume that every case of judging is like this one, that in every case criteria have to be decided on by the individual, and moreover that there is nothing at all that the individual needs to understand, either about the situation he confronts or about the language he shares with others, in order to be able to decide how, and therefore what, to commend.

In conclusion, then, Hare is taking for granted the common structure of shared language, and of shared responses to situations which that implies; and from within the security of that structure he puts foward a theory which seems to allow the individual to move outside it and judge from a wholly personal standpoint. He is ignoring what in a parallel discussion Wittgenstein called the stage-setting that needs to go on before such a process can occur. Inside the common structure of human language and the body of existing human judgement, dispute, discovery, discussion and revision of values can all take place; and this is what for the most part Hare is describing and analysing. But the structure presupposes a community of shared aims and shared responses which forms the background to all discussion of values and allows it to be meaningful. Where the values concerned are moral values, we can see the existence of the common structure underpinning our moral understanding when we realise that this or that non-prudential aim, and what counts as furthering or fulfilling such an aim, is bound to be relevant in coming to a moral decision on a particular question. For to understand that these are the moral aims people have is inseparable from an understanding of what 'good' means in a moral context. It is on this that I base my claim that prima facie moral judgements — the parameters, so to speak, of moral thought and decision —

*'Philosophical Investigations' ss242

are necessarily objective.

Now it might be suggested that this conclusion would be undermined by a different possibility from the one I have so far discussed. Granted that the individual cannot divorce himself wholly from the values of the culture he belongs to, and judge from outside it while still retaining an understanding of what value is, cannot we imagine wholly isolated societies arising in different places and having wholly different conceptions of value? This is, however, much more difficult to imagine than is sometimes thought. For it must be remembered that we are not discussing different views about what it would be <u>right</u> to do in a given situation, but something much more basic than that. In one society it might be thought right to kill or maim a poacher, and in another to do that might be thought quite wrong. Yet both these beliefs can co-exist within a moral standpoint which places value on freedom from bodily injury and also has a sense of the value of property. The difference here is one of decision about the circumstances in which the impulse of one of those values is outweighed by the other. In this respect, that of the relative weight of values in a given situation, it seems not improper to speak of a particular society's values — its moral code — as a consensus view. And it is partly this that the subjectivist has in mind when he urges the diversity of morals as an argument against the possibility of moral objectivity. It is also largely within this area that the moral reformer works. For often what he does is to urge the importance of one principle against the conventionally accepted claims of others that are in conflict with it, and so tries to alter our perspective within the moral standpoint. (To say that this is all the reformer does is to over-simplify matters, of course. For like anyone engaged in moral argument, he also points to facts which are unknown or unacknowledged, deduces consequences, points to inconsistencies in the conventional view, engages the imagination of others, and generally deploys all the features of rational moral argument that Toulmin, Hare, Nowell-Smith and many others have extensively analysed in recent years. He may also have the capacity to <u>extend</u> the moral viewpoint, leading the moral consciousness into

new areas of awareness, though this is necessarily rare.) Though here we may be discussing rival moral codes, then, it is not distinct and different conceptions of what morality is that we are being asked to consider. We must imagine instead a situation in which there could not be agreement at any level between members of the two separate cultures as to the goodness of anything.

Now Hare appears at first sight to be describing such a situation when he imagines an encounter between a missionary and a tribe of cannibals in 'The Language of Morals'.* Hare is again in pursuit of the notion that the meaning of 'good' lies in its use to commend, wholly irrespective of what is commended. He writes as follows:

"Let us suppose that a missionary, armed with a grammar book, lands on a cannibal island. The vocabulary of his grammar book gives him the equivalent, in the cannibals' language, of the English word 'good'. Let us suppose that, by a queer coincidence, the word is 'good'. And let us suppose also that it really is the equivalent — that it is, as the Oxford English Dictionary puts it, 'the most general adjective of commendation' in their language. If the missionary has mastered his vocabulary, he can, so long as he uses the word evaluatively and not descriptively, communicate with them about morals quite happily. They know that when he uses the word he is commending the person or object that he applies it to. The only thing they will find odd is that he applies it to such unexpected people, people who are meek and gentle and do not collect large quantities of scalps; whereas they themselves are accustomed to commend people who are bold and burly and collect more scalps than the average. But they and the missionary are under no misapprehension about the meaning, in the evaluative sense, of the word 'good'; it is the word one uses for commending. If they were under such a misapprehension, moral communication between them would be impossible.

We thus have a situation which would appear paradoxical to someone who thought that 'good' (either in English or in the cannibals' language) was a quality-word like 'red'. Even if the qualities in people which the missionary comm-

* op.cit. p.148

ended had nothing in common with the qualities which the cannibals commended, yet they would both know what the word 'good' meant..."

In denying that the situation Hare describes could ever come about if, as he suggests in the final sentence of the quotation above, the cannibals and the missionary have no moral beliefs at all in common, I do not mean to suggest that 'good' is a quality-word, or exactly like 'red'. But, like 'red', 'good' is a word whose intelligibility depends upon there being normal agreement on when it may properly be used, and hence on when a judgement that something is good is actually correct. Thus the ability of the cannibals and the missionary to discuss morals at all will depend upon their being able to agree, in a sufficiently wide range of cases, about what things are in fact good. This agreement need not be complete in order for them to be able to understand one another, and as I have indicated earlier in this chapter, the fact that 'good' changes its criteria of application depending on what it is applied to makes it very likely that agreement will be less than complete. For it unlikely that members of two such different societies will value quite the same things or value them in quite the same respects. Nevertheless, the area of agreement between cannibals and missionary must be extensive enough for it to be said that they would normally agree about a substantial range of moral cases. Hare's example, which appears to show that no agreement of any sort is necessary for understanding to occur, must therefore involve some epistemic slip or sleight of hand.

The first point we should notice about Hare's example is that the really difficult problem has already been solved before the missionary first appears on the scene. For someone has written a grammar book with a vocabulary which establishes that 'good', used as a moral term, means the same in both languages. And this is supposed to have occurred without there being <u>any</u> judgement which a member of one society would make or assent to, which a member of the other society would also make or assent to. But how has this feat of translation been performed? The grammarian has no doubt been able to recognise that when the

cannibals call something 'good' they are expressing favour or liking. But we must suppose that our grammarian, armed perhaps with a work on moral prescriptivism, has also noticed that the cannibals' use of 'good' in their language sometimes seems in accord with certain formal conditions: to wit, 'good' sometimes seems to be used to make universal prescriptions. But how has the grammarian been able to recognise that usage? Can he be sure that he has recognised it? And does the appearance of a universal prescription suffice to show the presence of moral concern? I do not think we can be sure of any of this.

Let us imagine a possible example. There are a number of streams running through the cannibal village, but only one of them is 'good' to drink from. Cannibal parents instruct their children to drink from no other, and no adult cannibal would himself do otherwise. We have something here that looks like a moral imperative, and might conceivably fool the grammarian. But if he is careful, he is likely to enquire into the reasons for this injunction. It might be because the stream is the only safe one to drink from, or again it might be because drinking from this stream brings the drinker luck. These are both prudential recommendations, then, and the grammarian would not on this account call 'good' in the cannibal's language a moral term. But now suppose instead that the reason for drinking from this one stream is the belief that to drink from any other will poison, not the drinker, but the rest of the tribe; or that drinking from the 'good' stream brings luck, not only on the drinker, but on the whole tribe as well. Will not the grammarian now be inclined to classify 'good' as a moral term? But there is no difference in the universal prescriptive uses of the term 'good' in these cases and the former ones. The difference is marked by the concern to do good to others, or not to bring harm on others. And this is something that the grammarian can understand, even if he doe not for one moment believe that any of these consequences will follow from drinking from one stream rather than another. For he recognises the appeal to beneficence as a characteristic distinguishing the moral from the non-moral case; and

this is a characteristic he is able to recognise as a moral one because it is one which is a moral one for him. For if he shared the cannibals' beliefs about the consequences to others of drinking from the stream, he would recognise those consequences as something which should guide his actions irrespective of the consequences to himself, or his own wishes in the matter. Now the remarkable thing about the example I have just given is that the grammarian will be quite unable to make use of reasons to distinguish prudential from non-prudential uses of the word 'good' in the cannibals' language if the cannibals' way of life is anything like Hare describes it. For if they value nothing that we value, the reasons that they give for calling something 'good' will never be anything that the grammarian could recognise as a reason to value that, and so he will never be able to distinguish the moral from the non-moral uses of 'good' in cases like those above. Indeed, when it comes to identifying the cannibals' fundamental moral principles, for which no reasons can of course be given, the grammarian will find no way of distinguishing a fundamental moral principle from a preference which is not felt to bind anybody except, obviously, the person who happens to have it.

Turning more specifically to the question of the relation between goodness and scalp-hunting, we can imagine a grammarian who believed in the doctrines of prescriptivism being misled in the following way into thinking that he had identified 'good' in the cannibals' language as a moral term. Suppose that the cannibals get a lot of enjoyment out of collecting scalps; and let us also suppose, for simplicity's sake, that they only use the word 'good' in their language in cases where they get this enjoyment. (This sort of personal satisfaction must be very common among them, if 'good' is to be a word of very general application. But that is possible.) Then we might imagine the cannibals all coming home in the evening after a busy day's scalp-hunting, and telling one another what a splendid time they had scalping today, and earnestly reinforcing their sense of enjoyment by recommending to each other that everyone should do it again tomorrow, because they

would all enjoy it so. Collecting scalps is 'good' — meaning that it's fun — and a 'good' man, the one who collects more scalps than others do, is to be admired and emulated because he must be getting so much more enjoyment out of life than others do. We can see why the grammarian might mistake 'good' for a moral term if it is used like this, because it is indistinguishable from a universal prescription, and he thinks that makes it a moral term. But suppose that the cannibals have no moral feelings one way or the other about scalp-hunting (in much the same way as I surmise people who go grouse-shooting have no moral feelings one way or the other about shooting grouse). It is just something they love to do. The cannibals are in fact using the word 'good' rather as children do when they describe their favourite activities as 'good fun'.

So the grammarian might be quite wrong if he thinks he has discovered a term of moral commendation in the cannibal word 'good' or even, strictly speaking, an evaluative term at all. For the word 'good' here might only express a sense of personal preference. While it is quite clear that the cannibals all favour scalp-hunting, it remains an open question whether they think it an activity which has any moral significance, or that a 'good' scalp-hunter is a virtuous man on account of his prowess in this direction. And so long as this does remain an open question, the grammarian cannot discover what the cannibals mean by 'good' just by the formal features of the word's use, and irrespective of what it is used to commend. The formal characteristics of moral language, which, according to the account given by prescriptivism, are the determining characteristics of language being moral language, just will not provide sufficient grounds for our being able to identify a term as a moral one.

Will any other characteristics which moral language or moral behaviour might have, irrespective of what it is about, suffice to identify something as a moral concern, and so enable the grammarian in this case to get to grips with what the cannibals mean by 'good'? It might be thought that this can be done by paying special attention to those activities which the cannibals take particularly

seriously, whatever they may happen to be. But there again, the cannibals might take scalp-hunting very seriously indeed, and think it very important to them, without it being the case that they derive anything more than personal satisfaction from it. They might think that others in the tribe would also get satisfaction from scalp-hunting, and so recommend the practice to them. The cannibal who refuses to go scalp-hunting might be thought strange to the extent that he becomes an object of scorn and derision — but yet there might be no question of this amounting to moral condemnation. In our society, chess players and car enthusiasts often take their activities extremely seriously, and may indeed proselytise vigorously. But the fact that they do this does not make chess-playing or enthusiasm for cars anything of any moral significance. It is not considered virtuous to play chess, for example, even among those whose lives revolve around chess-playing; and in a world where everyone else was a car enthusiast, the man who had absolutely no interest in anything to do with cars would not thereby be showing a lack of moral understanding, for all that he might seem a very odd person indeed to others.

Seriousness is a symptom, not a criterion of morality. But to turn in this direction in the search for a formal criterion of judgements being moral judgements may in fact be an attempt to formulate an abstract conception of moral content as a basis for identifying a society's moral concerns, whatever they may be. For, it is sometimes thought, if we look at the whole way of life of a society, surely whatever it is that its members regard as being essential for all will stand out clearly; and that will give us the key to their moral beliefs, and so enable us to identify their moral language. The notion that what people take seriously is a clue to what they value may be an inadequate attempt to express this idea.

To speak of what people think is essential to all is to speak of there being things that they think all people need. The idea that there are universal human needs is usually associated with moral descriptivism. But it might be imagined that if the grammarian can identify whatever

the cannibals think they need in order to live well or properly (in the sense that, while existence might be possible if those needs are unsatisfied, that existence cannot be a good one), he will have identified something which is logically linked to what they value. The trouble with this conception of how to identify the cannibals' values, however, is just that their needs are logically linked with their values. It is not that if the grammarian can identify their needs, he can identify their values. It is that he can only identify their needs (in the sense of identifying what they recognise as needs) if he has already identified their values. It is, we might think, quite possible to identify a good many human needs empirically. But that is only possible because there is complete agreement about the desirability of what these things are needed for. Even air, food and shelter are things we can be said to need rather than just things we want because there is complete agreement that we value survival, and these things are necessary for it. But can we make even that assumption about Hare's cannibals? Given that they value nothing that we value, it may be that they lead a sort of kamikaze existence, and care nothing for survival; or again it may be that, because nature is so kind to them that she provides for their every physical need without their having to make any effort, they are just not conscious of there being any value to things which they have never lacked, or imagined themselves lacking. It is still true, of course, that they do need these things for survival. But survival may just not have struck them as something valuable, so we cannot link the fact that they need these things in order to survive to the idea that they think them good. If the cannibals' view of life sounds a trifle odd, the oddity stems from Hare's original suggestion that it is possible for there to be people who value nothing that we value.

When we turn from considering those things which are essential if existence is to be maintained at all to those things which the cannibals think essential if life is to be a full and flourishing affair, it becomes even clearer that we cannot identify something as an ingredient thought essential to a good life by a cannibal until we are sure

about what the cannibals think a good life is like. And this forces us back to consideration of whether, when the cannibals say that scalp-hunting is good, they mean what we mean by 'good', or mean something different. Is scalp-hunting a need for them, or just something they like to do? The grammarian cannot be sure he has recognised scalp-hunting as something the cannibals see as essential to a flourishing life — essential, that is, in such a way as to be a moral matter for them, rather than just something very important to them as individuals but morally neutral so far as they are concerned — unless he has been able to do one of two things. Either he must have already identified their moral language (which, as I have already argued, he cannot do); or else he must have been able to assume that significant portions of the territory governed by their moral language is the same territory which his moral language covers for him (which, ex hypothesi, is a false assumption), and that there are really very many things which the cannibals value which he values too.

That this is a false assumption in the case before us is owing to the oddity of Hare's original claim. For in fact, even in the cannibal society as Hare describes it, there are possible points of contact between the values of the missionary and those which (pace Hare) the cannibals are likely to have. The cannibals, being "bold and burly", might be expected to recognise the virtue of courage, for instance, a concept which the missionary would also recognise. And in their dealings with other members of their own clans or families, at any rate, we should expect to find identifiable moral practices, which we would recognise, not because they were universalised or prescriptive, but because of whom they operated between and what they concerned. In short, we would recognise them as moral concerns because they are concerns which we share too.

I do not mean to imply by this that such qualities as courage, fidelity to one's comrades, respect for one's elders, or respect for property or the binding force of a promise, are to be taken as necessarily part of what is involved in standing within the moral outlook. Any or all of these qualities might be absent. In this respect I

differ from R.W.Beardsmore, who, in a critique of prescriptivism which is in some respects parallel to mine at this point, nevertheless suggests that a group of particular moral concepts including honesty, fidelity, truthfulness, justice and courage are necessarily present in anything that could count as morality, in that to be a moral consideration, reference must be made to one or other of these concepts.* I think it is highly likely that such concepts will be present in anything we would recognise as a moral matter, and very likely that they are present in some degree at least in a cannibal society as well. But there is nothing that dictates that they must be present in either one. Hare's recognition of the logical contingency of 'X is good' is to that extent correct; and to that extent it must also be a contingent matter that we find goodness in any of the particular virtues Beardsmore mentions. What is not contingent, however, is that our normal agreement on what we hold to be the content of morality is a precondition of our being able to use moral language intelligibly. Thus it can only be through the missionary and the grammarian finding positive values manifested in the cannibals' judgements where Europeans too would hold positive value to be present, that enables them to connect the meaning of the cannibal word 'good' with our word 'good'. The fact that missionary or grammarian are very likely to be able to do that only serves to emphasise how unrealistic Hare's claim that another society might have totally different values to ours is. For if we did find a group of people whose way of living involved nothing which we could treat as a point of contact with our own principles, so that there was never any behaviour which at some level of generality both we and they would commend, we would not find the idea that their behaviour indicated any moral concerns intelligible. For our entire moral vocabulary (including the word 'moral' itself) is intelligible only insofar as there is normal agreement in the conclusions we arrive at when we make judgements which concern these goods. Indeed, in the

* R.W.Beardsmore 'Moral Reasoning' (Routledge 1969)

absence of any such agreement between our moral parameters and the behaviour of a tribe we encountered, we might legitimately doubt whether we had encountered humans at all, even if these creatures were, biologically speaking, members of the same species as ourselves. For they would have nothing that could count as moral beliefs; and it is at least arguable that the possession of a moral outlook is a necessary condition of being human.

In defence of the prescriptivist's notion that the intelligibility of a moral judgement is something quite distinct from the question of whether anyone else agrees with it, and in part as a reaction to Beardsmore's argument to the contrary, W.D.Hudson has presented an ingenious counter-argument in the form of an analogy.* In this, moral evaluation is compared with pricing, which is after all a form of valuation. Comparing the world in which moral judgements are made with a market place in which things are put up for sale, Hudson suggests that if everyone in the market puts a different value on the goods he is trying to sell from the value anyone else would put on those goods, so that no-one ever <u>does</u> agree on the price of anything, the act of pricing nevertheless remains an intelligible one. Hudson is careful to make it a contingent matter that no bargains are ever struck in this market, and not a logical matter that disagreement is endemic, so that no bargains <u>could</u> ever be struck. But his position is still, I think, an untenable one.

Like so many analogies, Hudson's trades on the familiarity of the situation which it describes. For it seems unexceptionable that someone should put a price on his goods, but that in everyone else's eyes that price just happens to be <u>wrong</u>. We know what a wrong price is because we know what a right price is, and so it is conceivable that everybody might happen to set prices which everybody else finds wrong. Now this situation makes sense in a market where everybody agrees that all the goods on sale have <u>a</u> price which is right, but just cannot agree on what that price is. (Thus the analogy is a reasonably

* W.D.Hudson 'Modern Moral Philosophy' (Macmillan 1970) p.311

accurate representation of something that might happen within the moral standpoint, where everyone agrees on what has value, but no-one happens to agree with anyone else on the relative value of anything.)

But that is not the situation which the analogy would have to represent in order to counter the arguments I have been putting foward against Hare. For the man or the society that is supposed to be operating outside the moral standpoint which we all share is not disputing the relative value of things which are agreed by everyone to have some value, but asserting that something has value when nobody else thinks it has any. Now it looks like an act of pricing when somebody names a figure while indicating something, but this is because we are all used to the idea that people only sell in circumstances where others are likely to buy because there is agreement that the things in question do have value. And talk of buying, of selling, of right or wrong prices, and of pricing at all, makes sense only in that context. But the man who tries to sell something which only he thinks of as having a value at all is not operating within that context. It is not that he puts the wrong price on his goods, but that he is wrong to put a price — any price at all — on them. This becomes evident if we imagine a man trying to sell something which it would not occur to anyone else to think of as having a price (because, say, it would not occur to anyone else that here was something that could be bought or owned or transferred). Thus the man who tries to sell the visible spectrum, or the number seven, or his own aesthetic appreciation of Mozart's Symphony no. 40, or Antony's love for Cleopatra, does not fail to find a buyer because he has misjudged the market for such things, and set his prices too high. For there is no market for such things, and they do not have prices. In the same way, the man who proposes that the action of not stepping on the cracks between the paving stones has more intrinsic moral worth than the action of keeping one's promises, has not made an error as to the comparative moral worth of these two types of action. He has made an error of a different sort

in thinking that they can be compared at all in those terms. For only one of these types of action has any intrinsic moral worth at all.

Hudson's analogy fails, then, because in order for it to be an apt analogy it must presuppose the very agreement about what things have value which it sought to deny the need for. Thus it lends no support to the view that a person or a society could frame a personal conception of what counts as the morally good, and at the same time maintain an intelligible claim to understand what moral goodness is. But this is inevitable. For, in the first place, as I argued in the early part of Chapter Five, following Wittgenstein, any attempt to divorce our capacity to understand a concept from the existence of a set of public rules dictating in what circumstances the concept may be meaningfully used in judgement, is bound to fail. For meaningfulness demands such public rules. This by itself would be enough to outlaw the usage of 'good' invented for himself by the person who claims to occupy a position outside the moral standpoint. For his usage, being governed by his own impressions of self-imposed rules, can be neither right nor wrong.

But according to my interpretation, it will be remembered, Wittgenstein is going further than that. He maintains that where it makes sense to say that a judgement employing a concept can be right or wrong, there must be normal agreement in the results of our judging ("a certain constancy in results", as he calls it*) if such concepts are to be used meaningfully to communicate. Thus any concept which can occur in a judgement which may be the subject of agreement or disagreement, and where justification by reference to evidence is possible, is a concept which can only have meaning for us if we not only agree when it has been used meaningfully, but also agree when, in general, judgements containing that concept are actually right. Goodness is such a concept, for we can and do dispute about what things are good, and bring foward reasons and evidence in support of our judgements;

* 'Philosophical Investigations' ss242: (see p.126ff. above)

and from this it follows that our understanding of goodness requires us normally to agree about what things are in fact good. This does not entail that 'X is good' is a tautology or entail that we must all always agree in every judgement we make about what is good or again entail that our understanding of goodness resides in our capacity merely to judge as others do. For the possibility of meaningful disagreement still exists, and has a multitude of sources. Yet given a normal understanding of the circumstances in which it is right to call something good, it should in principle always be possible for one party to a dispute to show another that he has judged erroneously, either because he has made a mistake about the evidence or an error in reasoning, or else because he has not properly grasped what it means to be good in the context under discussion. It may not always be possible for someone to get his opponent to see that something is good. Yet the standards enshrined in the way people normally judge are the standards which must determine the matter, whether or not a man can appreciate what they are in a particular case. For these are the standards from which any conception of what constitutes relevant or decisive evidence must flow, since these standards embody the criteria for correct judgement. And where the judgement in question is a fundamental one, so that considerations of criteria and evidence follow from it rather than determine it, it can only be the public agreement that this or that counts as a correct judgement of its kind, expressed by the fact that this is how people normally judge in this matter, which makes it possible for the concept employed in the judgement to find meaningful expression in language, and so be something which can be communicated to others and be thought by oneself.

This general argument for the necessity of there being normal agreement in fundamental judgements if the concepts employed in the judgements are to be meaningful, which I charted in Chapter Five, shows it to be inevitable that our most basic judgements about what things are morally good should be objective. For the very meaningfulness of the

term 'morally good' depends upon there being this agreement in our judgements about what things, fundamentally, *are* morally good. And that is the conclusion we have also come to, by a slightly more pragmatic route, in the course of the last two chapters. For in considering the contrary position, the belief that anything might be called morally good, and for whatever reasons we choose, we have seen first of all that without agreement in what counts as evidence, both the possibility of genuine moral dispute and the possibility of its resolution by rational argument disappear. More conclusively still, we have seen that without there being normal agreement in what in fact is good, moral concepts themselves become unintelligible. This constitutes a *reductio ad absurdum* of the notion that what counts as moral goodness could, in the last analysis, be something about which people could consistently yet meaningfully differ; and that in itself guarantees that when a man judges that some state of affairs or some action has moral worth, his judgement will be either right or wrong according to a single standard which is genuinely independent of his own or his society's beliefs, and yet which his beliefs and those of his society will normally reflect as far as their understanding permits. To see this is again to see that, to the extent that we judge with understanding, and in the normal way, our judgements about what is morally good are necessarily objective.

CHAPTER NINE

MORAL OBJECTIVITY AND MORAL DIVERSITY

The programme I have set for myself in this thesis is now complete. I have explained how empirical judgements come to be objective, by showing that our judgements in this field can be rational, and the concepts we employ can be meaningful, only if there is general agreement among us in the results of our most basic empirical judgements. And, I have argued, the same is true of moral judgements. Our moral discourse can only be rational, and the concepts we employ will only be meaningful, if there is general agreement among us in the results of our most basic moral judgements, the judgements about what things are good. Thus these judgements too are objective when we judge with a normal understanding of what we are doing, and so in consequence are all those other moral judgements which can be derived from them. However, this conclusion is by no means as sweeping as it might appear at first sight. So in order to avoid misunderstandings and objections which might arise from them, I want to conclude my thesis with a brief discussion of some of the limitations in the scope of moral objectivity, as well as saying something about its importance.

A standard, though misguided, argument against there being any form of moral objectivity is the appeal to the variety of forms of behaviour which have been held, at different times and in different places, to be morally right. That different people or different societies think different things right is in itself no more a disproof of moral objectivity than the fact that children get different answers to their sums is a proof of the subjectivity of mathematical calculation; and many of these differences may be discounted when we realise that the form moral behaviour takes may vary widely on account of circumstances and beliefs which are only contingently related to the moral dimension of the behaviour involved, but which may strike an observer as simply wicked because of their

divergence from his own customs and practices. This may be because he just does not understand the moral import of such behaviour, or else because he could not bring himself to express his beliefs in actions of that sort. Into this category of discountable differences fall such things as the ways in which different tribes show respect for their dead, which Herodotus remarked on. But differences of this sort should not prevent us from treating respect for the dead, which is an integral part of respect for persons (given that a person's body is thought to be part of his person, or at least an honoured possession of his), as something which any moral man should have a concern for. And when we have taken account of some of the other sources of the diversity of moral judgements which I shall indicate in the next few pages, this diversity of forms of expression of moral beliefs, as it might be called, will be seen as having no great philosophical significance.

Rather more important, I believe, for our understanding of the sources (other than error) of the diversity of moral judgements are two limitations inherent in the theory of necessary agreement as the basis for objectivity, in the form in which I have put it foward. The first of these limiting features is simply that our understanding of what is objectively good may well be incomplete. (This would be the case, I would argue, in a society which showed respect for persons but none for their dead bodies, for instance.) I have maintained that there must be a body of beliefs about what is good which is common to all of us if we are to understand one another's moral discourse at all. But it does not follow that your beliefs and mine must be entirely co-extensive. Where I lack beliefs which you have, or where I have beliefs which you lack, you will be entitled, insofar as you are a normal judge of these matters, to regard me as lacking understanding in some measure with regard to morals — and conversely, I might regard you as deficient in moral understanding. It is then open to either of us to try to improve the understanding of the other by a process similar to ways of getting someone to grasp other sorts of concept more fully or more

accurately, such as getting the other to explore the possibilities of the application of the concept in situations which he has not considered before, or suggesting that its application in certain circumstances really involves a confusion with some other notion. Experience, imagination and reflection all play a part in this process; for we cannot be said to have a grasp of the applicability or otherwise of a concept to a situation without cultivating some idea of what it would be like to be in that situation. This is not an appeal to the significance of the feelings which may or may not accompany such actual or imagined experience, and may sometimes act as a prompt to, or even at times a substitute for, moral insight. It is rather a reference to the epistemological point that understanding requires there to be exposure to experience in order to be anything more than formal understanding. Thus, and most commonly, a person's failure to see the moral worth or otherwise inherent in a particular situation may stem from the fact that he has simply never encountered or considered anything like it before, and has no idea of what to make of it.

Furthermore, even if all the members of a society had exactly the same moral understanding as one another, there is another limitation embodied in my view of objectivity which allows disagreements in judgements to occur at various points — though not, it should be said, to any greater extent than might occur in any other form of judgement that can be objective. I am not thinking here of our ordinary liability to be mistaken about the facts of the case, or to make invalid inferences, which again are to be found in other types of judgement where objectivity is possible. I have in mind the fact that our mastery of a concept requires only that we should <u>normally</u> agree with others in our fundamental judgements, and not that we <u>invariably</u> should. It is thus possible that I should sometimes simply fail to recognise the goodness of a particular situation, even though I normally would, or that I should see moral value in something that just does not possess it, though again I normally would not do this. The explanation of such aberrations presents great problems.

That they can occur may be just a brute fact about us, like the fact that our memories sometimes play us false. Alternatively, these aberrant judgements may be explicable by reference to logical factors associated with the existence of the complex "psychological penumbra" (as G.J.Warnock has called it) of feelings and emotions which surround the process of moral judgement. When a moral judgement is made, there may be nothing present which would, strictly speaking, count as an inner feeling at all. But again, there may be. And where such inner feelings do occur, in an individual, they may come to assume a place as an integral element in the process of moral judgement for him. This is at least conceivable. Now if such feelings do occur, in a manner analogous to the occurrence of inner experiences associated with sensation, they must be defined and identified by reference to what, objectively and publicly, counts as the morally valuable. That is, our capacity to identify such experiences as what they are for us would have to turn upon public agreement in the judgements involved. But then, in the same way as the inner experience in perception comes to have a complementary but distinguishable role from the public agreement in judgements which allows us to recognise the inner experience for what it is, and lets us judge on the basis of the inner experience itself, so the role of these inner 'moral sensations' may achieve or approach the status of an inner criterion for moral worth, at least for some of us. Someone's having, or lacking, a particular feeling in connection with some state of affairs may then be grounds, for him, for saying that the state of affairs before him has, or lacks, moral worth. And if this is wrong, and his judgement is consequently abnormal, at least its subjectivity is explained. The relation of our feelings and emotions to our moral understandings raises interesting and difficult questions, and seems to be the likeliest source of insight into the point at issue here. But I cannot go into these questions in the present work, beyond offering the thought that the form my theory has

* G.J.Warnock 'Contemporary Moral Philosophy' (London 1967) p.52

taken does allow for a distinct but complementary role to be given to moral feelings, in some respects parallel, but not identical, to the role of reason in moral judgement.

The sources of moral diversity which I have discussed so far are by no means trivial, but they do not reach to the heart of the matter. If all moral questions could be settled, as a rule, by reference to our agreement about what, fundamentally, has moral value, the diversity of sincerely held beliefs, even within our own society, would have to be explained as errors due to faulty or incomplete reasoning, false or restricted data, or widespread lack of moral understanding, or a combination of all of these. And while part of this is true in a certain sense, this is not something we can lay solely to the charge of one another as individuals. For our capacity to be objective about moral goodness is (so to speak) the touchstone for the justification of moral beliefs, rather than the determinant of all our moral judgements. Moral diversity is a consequence of characteristics inherent in the structure of moral thought itself — or, to put it another way, a consequence of inherent limitations in the scope of moral objectivity. These limitations are reflected in the structure of our moral discourse, and in the way that those of our moral concepts that relate to conduct function, as I shall now try to explain.

The position I have been developing in this thesis is very much in contrast to most traditional objectivist theories of morals, not only because the type of justification I have based my case on is markedly different from traditional appeals to the objectivity of the individual's moral intuitions, but also because I have not been claiming that to every substantive moral question there is an objective answer available to us. I have only said that our fundamental judgements as to what is morally good will normally be objective, because this must be so if people are to have a communicable and intelligible sense of moral value. Now it will be the case that any other judgement which follows, either as a matter of logic, or else on the evidence, from one of these fundamental judgements will

also be objective. But many of the moral judgements we make will not qualify in that way as objective. In particular, judgements about what is morally right or obligatory do not have any automatic claim to be taken as objective, and I have been careful not to suggest that they have. In order to put my advocacy of the objectivity of moral goodness into perspective, then, I want to indicate briefly why I think it would be a mistake to treat judgements about what it is right to do, or about what we ought to do, as objective in their own right, or indeed, for the most part, as objective at all.

Now it might be thought that our judgements about what is right or obligatory, in their most basic forms, would qualify as objective in just the same way, and for just the same reasons, as the fundamental judgements about moral goodness which I have been discussing. For, in the light of the theory I have put forward, it looks as if fundamental judgements about what we should do cannot admit of any further justification apart from our normal agreement that these terms are to be used in certain ways to make certain judgements. But although many philosophers have regarded such judgements as sui generis (and therefore, if objective, then objective in their own right), I believe that in doing so they have misrepresented the role such judgements play in moral discourse and moral thought, and failed to see how they arise in response to the practical exigencies of moral decision-making. Where judgements about what we should do are objective, their objectivity depends finally on their relationship to what is objectively good. But for the most part, such judgements reflect our inability in practice to attain the conditions for objective judgement. The interrelations between the terms 'good', 'right', 'the right', 'ought' and 'permissible' are extremely complex, and involve many subtle distinctions which I cannot go into now. There is, I think, valuable work still to be done on the conditions under which a judgement involving one of these terms justifies or is justified by a judgement involving another of them, in the light of my argument that only fundamental judgements involving one of these terms, and those other judgements which can be derived from them,

are in fact objective. What I can do here, however, is to sketch in those general considerations which lead me to think that judgements about what we should do in moral matters are not objective in their own right, that where they are objective, their objectivity is due to some objective good to which they directly or indirectly refer, and finally that they are for the most part not objective at all.

The real basis of the case for saying that judgements about what is right or obligatory are independent of other types of moral judgement (and that therefore in their fundamental forms their justification must depend solely on our general agreement on how we shall judge in such matters) is that these judgements seem to be independent of considerations of the consequences of the actions they dictate. Thus the judgement that X is the right thing to do in a situation of a certain kind does not forfeit its claim to correctness even if, in a particular case, X turns out to produce more harm than good, or even nothing but harm. The simplest justification for saying that X is the right thing to do, even in this case, consists in pointing out that the doing of X conforms to a <u>rule of conduct</u> — that, given the kind of case involved, X just is always the right thing to do. This justification relies on our natural respect for rules of conduct. But our respect must have some source; and a more sophisticated explanation justifies the value of acting according to a rule in terms of the point of doing so.

Part of the point of employing rules of conduct to govern our judgements is that they facilitate the making of such judgements. For they enable us to formalise methods of bringing about desired ends. Doing X may cause harm in certain cases, but if in the majority of instances it guides us to do what produces good, it is obviously a useful rule, particularly as in many cases we have no better method of deciding what to do, a point I shall have more to say about later. I do not mean here that the justification for following a rule is simply that in most cases doing so will result in the production of good, so that we could provide a translation of every judgement which a

deontologist would regard as *sui generis* into one which refers to the production of some good, with the addition of the word 'probably' to cover those cases where following a rule fails to lead to the appropriate consequences. Things are not quite as straightfoward as that! Yet we should notice that although an action may be right or obligatory even where it brings about harm, we cannot be indifferent to that harm, even though we feel that the harm done in particular cases is overridden by the action's being the right thing to do as a matter of principle. And if it turned out that action in accordance with a given rule *never* produced anything but harm, we would be quite entitled to doubt whether that rule *did* in fact tell us what the right thing to do would be. This consideration indicates, if it does not establish, the dependence of judgements about the right and the obligatory on judgements about what is good.

The dependence of judgements made in accordance with rules of conduct upon judgements about the good ends to be attained must be established by rather more general considerations, however. For it is the *general* relationship between judgements concerning duties and judgements concerning ends which shows the *general* dependence of the former upon the latter. I have just made the point that a complete absence of good consequences would constitute a powerful rebuttal of a claim that it is right to act in a certain way. But we can imagine someone clinging to his faith in a particular rule of conduct in the face of this rebuttal, and we would have some sympathy for his position if the rule he was clinging to was not an isolated one, but an integral part of a larger system of principles of conduct. Yet when we *generalise* from the absence of concern for good consequences in relation to a *particular* rule of conduct, to the situation that would obtain if rules of conduct *as a whole* took no account of what happens if people follow them, the position rapidly becomes an untenable one. It would be odd indeed to suggest that there are all sorts of things which it is right or obligatory to do, but at the same time to deny that there are *any* good ends that such actions could be expected to bring about, or to deny that

what ends are brought about is a matter of any relevance at all. It would be odd to think this because the language of proper actions is logically dependent upon the language of worthwhile ends as part of the logical dependence of the language of actions as a whole upon the language of preferred ends, or, as I have called them in earlier chapters, aims or goals. The latter forms a background without which the former would not be intelligible. For it makes no sense to talk of actions at all unless we have an interest in things being one way rather than another — that is, unless we have preferences. Thus the general idea that there are things it is right to do presupposes the general idea that there are states of affairs that it is good to bring about. But the converse is not the case, as we can see in the context of morals if we reflect on the absurdity of the idea that good states of affairs might be good only because right actions produce them. It would be very odd to believe, for example, that freedom from pain is a desirable end only because it is right to help the suffering. But it is perfectly natural and correct to believe that what allows us to talk of actions in general as being morally right or wrong is the fact that we are not indifferent to how things are, that action is the means of bringing about changes in how things are, and that there are some states of affairs which merit our trying to bring them about for their own sakes. Thus, while we may not be able to translate each and every judgement about what is right or obligatory directly into a judgement about a good to be attained, we must recognise nevertheless that our talk of morally proper types of behaviour gets its point from there being various types of situation that we want to bring about. It is in this sense that judgements about what is right or obligatory must depend for their ultimate justification on our views about what is morally good. Hence any objectivity pertaining to judgements about moral conduct must in the last resort depend on the objectivity of our judgements about what is morally good.

We see something of this dependence of conduct-language on the language of good ends in one use of the word 'right' where we can sometimes claim objectivity for the judgement

that a particular action is right. This is a partly permissive rather than a wholly directive use of 'right', where we are not saying that something is the right thing to do, but that it is an action which may rightly be done. Yet at the same time we are saying more than that it is merely 'all right' to do this — more, that is, than saying that the action is a morally neutral matter which no rule of conduct either enjoins or prohibits. For we would commend a man for doing a right thing, at any rate insofar as this involved no conflict with doing some other thing which in these circumstances was the right thing to do. To call an action right in this sense must be to commend it for being, directly or indirectly, conducive to a good end, as I have argued. And the most obvious way in which this could occur is when we call an action a right one, not because it conforms to a rule, but simply because it is directly conducive to the production of some good and we have no reason to believe that acting in this way will cause harm in any other respect. It may be far from easy to be sure this second condition is satisfied in practice, a problem I shall turn to in a moment. But if we can be sure of this, we can say that the judgement that this is a right action is an objective judgement just because of its relation to the good it produces. Such a judgement, though objective, is only slightly more informative than the corresponding judgement about goodness from which it gets its direct justification. The judgement that such-and-such is a right action does not settle the question of what we should do if there are any other actions which would also be right to do in these circumstances because they conduce to other goods. And the impulse to do a right deed need not be seen as coming from the rightness of the action or from the force of a moral law under which it is a right deed, but can come instead from our recognition of the good involved, which constitutes an aim for us just by virtue of the fact that we recognise that it is a good. The objective rightness of an action, in this sense of 'right', is thus so completely subservient to the goodness promoted by so acting that we are entitled to regard 'right' here as merely what might be called the performative form of 'good'.

To act rightly is, in this sense of 'right', just to act well. And it was on account of this that I spoke in some previous chapters of there being morally good actions as well as morally good states of affairs. For it is perfectly natural to extend the goodness ascribed to states of affairs to the actions that bring those states about, just because they do so.

But if objectivity (albeit derived from the objectivity of judgements about the good) can fairly readily be a feature of judgements about the rightness of those actions which are directly related to the goods they bring about, it is very much less likely to be a feature of judgements about what, in a given set of circumstances, one ought to do, or what would be the right thing to do — the central classes of moral judgement relating to conduct. For here there is so much more that has to be taken into account; and much of what would need to be taken into account in order for our judgements to be objective is simply inaccessible to us in practice. It is for this reason, I suspect, that rules of conduct have the moral significance that they have for us. For these rules are usually venerable pieces of public property, serving not only to guide our own conduct but also to regulate the conduct of others whom we advise or influence; and as such they represent a valuable repository of collective experience and judgement. Furthermore, these rules constitute a source of ready-made 'in principle' judgement, which we can bring to bear on the actual cases which confront us, as guiding if not determining principles, even though we may be well aware that judgement in principle is often artificial to the extent that it represents moral dilemmas in simplified form, and relies on the use of ceteris paribus clauses in order to render its conclusion unequivocal. But this is the price we must pay for a method which allows us to dispense to some extent with having to decide every case on its own merits, in relation to the good that would actually be produced by the actions we contemplate — or rather, with having to try to decide this. For in order to decide, on this basis, what actions we ought to perform, we would have to take all morally relevant considerations

into account, and be sure we had left none out of the reckoning; and we would have to know what all the consequences would be of every course of action which it lay in our power to perform in the circumstances. And clearly this is a counsel of perfection. It was the recognition of this fact which led H.A.Prichard to argue that a person's obligations are determined in what is essentially a subjective manner.* Recognising that obligations belong to individual agents rather than to actions themselves, and recognising also the absurdity of making an account of obligation require the agent's complete knowledge of his situation and all the consequences of all the actions it would be possible for him to perform in that situation, Prichard maintained that a man's actual duty is to do what he thinks right in the situation as he thinks it to be. It cannot be his duty to do what is actually right in the situation as it actually is, for the practical impossibility of his knowing all this entails that an agent could never, in practice, know what his duty was, and might do, or fail to do, his duty without knowing that he was doing, or failing to do, it.

Prichard's argument presents certain difficulties, not least that of accommodating the distinction between its being my duty to do something, and my only thinking that it is. To overcome this difficulty, we need to recognise that the agent is capable of having reservations about his own beliefs, if not at the time he makes his judgement, then later on; and that other people are also entitled to cast doubt upon whether the agent's beliefs, at a particular time, constitute a suitable basis for him to make a judgement at all, even though it must be the agent's own beliefs which in the last resort will determine what he ought to do. But having voiced these reservations, it must also be said that Prichard's central insight is sound. He is certainly right to say that an account of duty which makes it impossible in practice ever to know what one ought to do is absurd. Thus at some point we must accept that it is an agent's beliefs — reasonably well founded, perhaps, but

* H.A.Prichard 'Duty and Ignorance of Fact' (1932) reprinted in his 'Moral Obligation' (Oxford 1949)

nevertheless his beliefs — which determine what he ought to do, and not what is really right in the circumstances as they really are. For this alternative is simply untenable.

From this it follows that if two men see their obligations differently, although the situation is in fact the same for both of them, neither of them need be unjustified in his claim to have judged rightly in deciding what his obligation is. For their obligations will be as they see them, at any rate providing that they have made their judgements with reasonable care. And the supposal that judgements about obligations cannot be justified unless they are objective (so that we all either agree in our judgements or else are mistaken about what our duty is) is simply false. A man's judgement about his duty will only be objective insofar as his beliefs about the circumstances and consequences on which he based his judgement are actually complete and correct. And it is precisely because this is so improbable that we have a use for a concept of duty which is justified on the subjective grounds which, in broad terms, Prichard's argument suggests.

The situation with respect to the objectivity of judgements about the right thing to do is rather similar. For, as I have argued, the question of whether an action is right must at some stage depend on questions about the production of goodness by the proposed action, or by actions of that sort performed consistently in accordance with a rule of conduct. Thus there is further scope for the argument that our judgements about the right thing for us to do must be grounded in our beliefs about what will be productive of goodness, rather than in knowledge of what will actually do so. Once again, we are faced with the problem that we often do not know how things will turn out if we act in one way rather than another, and we often cannot be sure that we have taken proper account of every aspect of the situation that the judgement is a response to. So in order to avoid the sceptical conclusion that we can never be justified in saying that such-and-such is the right thing to do, we must hold that our beliefs about the general consequences of our acting in that way are

what justify us in saying that this is the right thing to do. Once again, such beliefs must be relatively well founded in order to allow for the distinction to be made between something's being right and its only being thought right. Yet there must come a point where we say that an action is right, on the strength of our beliefs about what will come of it. It is a paradox of action that if we never act in a case unless we are completely sure of what the outcome will be, there will be many departments of life where we will never act at all. Morality is one such department — arguably the most important one. So we must expect to find that the right thing for one man to do, in a given situation (and given his beliefs about that situation and the consequences of actions), is not necessarily the right thing for another man to do in the same situation (and given his beliefs) — and it is not the case that at least one of them must be wrong in his moral judgement. At least one of the men must have a less than total grasp of the circumstances of the case, or a less than complete understanding of what the consequences of all the possible courses of action open to him would be. (And of course it is also possible that at least one of the men has a less than adequate understanding of what things are good, or is suffering a temporary lapse from his normal ability to judge what is good, as I explained earlier.) But doubts about the strength of the evidence, at any rate, need not be translated into reservations about the soundness of the moral judgements themselves as judgements made on the evidence available, and made with reasonable competence. Thus we can accept that the correctness of people's judgements about what it is right to do need not entail that they should all reach the same conclusions about what is right.

The scope for diversity among moral judgements is even further extended when we take into account another feature which is a characteristic of the dilemmas which often confront us. So far in this chapter I have tended to speak of the actual production of goodness, or of a good state of affairs, as the objective standard against which judgements about conduct must finally be measured, to the best of our ability. This was convenient, and did not affect any of the

arguments put forward there; but nevertheless, it was an oversimplification, and one with significant consequences for any remaining possibility of there being objective judgements about moral conduct. For our aim in trying to decide how we ought to act is not to bring about something abstract but unitary — something called 'the good' — but to bring about any or all of a number of different and relatively specific good states of affairs. Difficult moral problems are usually so not only because it is hard to get to grips with the evidence which would suffice for a sound decision, as I have argued above, but also because we frequently find ourselves in situations where we are faced with dilemmas about which good to pursue, of the many which we might pursue in those circumstances. Often we must decide between competing goods, and decide in the knowledge that the attainment of one good end is likely to be at the expense of not attaining another, or at the cost of actually producing harm in some other respect. Our realisation that the goals we must decide between are all objectively good will not help us here, when our circumstances will not permit us to do all and only those things which lead to good ends. And it was for this reason that I referred in earlier chapters to objectively good states of affairs as prima facie goods, not implying by this that they might turn out not to be good after all, but simply to guard against the idea that our recognition of some good as a good means that we ought to pursue it regardless of the consequences that might have on the pursuit of other goods.

Now it may well be that, for us, some goods clearly and even normally take precedence over others. The value of preserving life is obviously a candidate for the status of prime good in many cultures, though equally obviously it is not regarded as an indefeasible good, since every culture seems to accept that there are circumstances where actions which result in the loss of life would be morally right. These are not necessarily the same circumstances, though; and this consideration alone would be enough to suggest that our conception of the hierarchy of goods is at least partly a matter of perspective. Perspectives can

change, and with them the moral codes that give them expression. And we cannot assume that other individuals or other societies will necessarily share our perspective, and so share our system of moral priorities. Nor can I see that any particular moral perspective should necessarily be more correct than any other, providing always that all these perspectives encompass the full range of moral goods. With the recognition of the possibility of there being different perspectives on the relative value of different moral goods we witness the disappearance of any remaining possibility of objectivity in relation to judgements about moral conduct. For the relative value we accord to different goods can be a decisive factor influencing the answers we give to the question of how we ought to act in a given set of circumstances. We cannot maintain that, given complete agreement among men in what they believe the circumstances to be, and given complete agreement too in their beliefs about the consequences of the various courses of action open to them, there will be universal agreement on what must appear to all as the right course of action. Thus if full objectivity in judgements about conduct is impracticable so too what I have earlier called objective apparent judgement would seem to be unattainable in our judgements about moral conduct. For, all other things being equal, our perspective on the relative importance of different goods may still lead us to judge differently from those whose moral perspective is different from ours.

It may be thought that the view I have sketched in over the course of the last few pages offers an unduly sceptical assessment of the scope for objectivity among judgements about moral conduct. It does not seem to me to be an unrealistic view, given the diversity of sincere moral beliefs which we find among men. But I accept that the question cannot be regarded as settled by what I have said here. However, it may also be thought that my scepticism about objectivity with regard to judgements about conduct tends to undermine the importance of my main conclusion about the objectivity of judgements as to what is morally good. Yet the very opposite is in fact the case.

For the more sceptical we are about the objective possibilities of other forms of moral judgement, the more significant the objectivity of judgements about the morally good becomes. The more we see our moral decisions about conduct as relative to our own beliefs and our own perspective within the moral standpoint, the clearer it is that the moral standpoint, formed as it is from our normal agreement as to what is good, carries upon itself the whole task of identifying and underpinning whatever is distinctly moral in our thoughts and actions. For if our normal agreement as to what is morally good is the sole repository of moral objectivity, it must be this that gives the term 'moral' its distinct identity as a concept, without which moral discourse becomes indistinguishable from the mere expression of personal preferences, except that the latter are more honest in that they are not disguised in a special language to give them a spurious air of dignity and independence. In addition to this, the more we see judgements about what is morally good as the sole source of moral objectivity, the greater our reliance on these judgements must be as the determinants of the content of morality. And the idea that morality has a determinate content is an essential one. It both establishes a series of definite targets for our conduct to aim at, and sets objective limits to what could count for anyone as a moral decision or a moral consideration. Without determinate aims and definite limits, morality becomes quite literally anything we choose, so long as we dress what we choose up in the garb of personal consistency. And moral dispute and argument is reduced to something in which, in C.D.Broad's phrase, all we can finally do is 'to twit each other with inconsistency'.

Thus, in the absence of any other sources of objectivity in morality as a whole, reference to actual moral goods becomes the touchstone of both moral relevance and moral justification. Such questions as 'but will this, after all, lead to some good being attained?' and 'but how is this related to the attainment of some good?' become the test of moral purpose, of moral relevance, of reasonableness and of correctness within all moral argument and all moral

thought. It may not be the case that we can determine, just from our knowledge of what is in fact good, what we ought to do in any particular situation which calls for a moral response; yet it must be the case that our capacity to recognise the need for a moral response, and our capacity to recognise an answer as being within the permissible limits of moral response, have their origins in our knowledge of what is morally good. The test of whether an action is conducive to the objectively good gives to the process of moral decision, of whatever kind, and at whatever level, a content, a yardstick and a goal. None of these are ours to determine for ourselves as individuals. For without their independence we could not recognise morality at all as indicative, in the final analysis, of anything beyond the generalised expression of our own particular interests and desires.

—oOo—

APPENDIX: IMMANUEL KANT

Kant maintained that both empirical and moral judgements can be objective, and discerned an important degree of similarity between the reasons for that objectivity in either case. In both these respects I have followed him, and my work owes much to his influence. To do justice to Kant's thought on the nature of objectivity would, however, have required a very different work from the one I have written. I would have found it hard, if not impossible to avoid letting the range, the direction and the manner of his thinking dominate my own; and for the reason given below I felt it necessary to develop my own approach to the problems I discuss in this thesis.

Taking my own line has resulted in considerable differences between my final position in this work and Kant's own conclusions. I am obviously at variance with him over the location of the source of moral objectivity, as well as over the extent to which it can govern our judgements. Kant locates the foundations of objective moral judgements in absolute moral law which we recognise it as our duty to be bound by, whereas moral objectivity's humbler abode in my account is a number of good states of affairs which we must agree to recognise as having value for us all, and which guide our conduct in a rather more modest way. I have given some indications in my final chapter of why I think the worth of states of affairs is logically prior to the worth of actions performed as moral duties. And with this would go the belief that the Categorical Imperative could only be a test of whether a principle is a moral rather than a prudential one, and not a source of objective moral law when combined with the exercise of reason. For, as I argue there, I do not think that reason can in practice ever carry us to a position where we can know, without doubt or equivocation, what we ought, absolutely, to do.

However, the origins of these differences lie in something which begins to show itself very early on in this thesis, at the point where I start to suggest that objectivity is distinct from truth. For Kant believed

that if a judgement is objective, it is necessarily true. This is a consequence of Kant's whole method of constructing the Critical Philosophy; but it results in an account of objectivity which I find seriously defective, if only because Kant's theory merely accounts for the objectivity of those judgements where disputes never arise — judgements which are necessarily true. In order to explain why I believe this leaves out most of what is important and interesting about the concept of objectivity and its applications, I will very briefly summarise the relevant aspects of Kant's theory.

One of the problems posed for Kant by Humean scepticism was that of giving an adequate account of our knowledge of the existence of an external, independent world, one that is the same for all of us. This embraces the question of our knowledge of objects (though it embraces much else besides). Kant regarded our sense-perceptions as subjective intuitions only, merely held together by our capacity for perceptual judging. These perceptions can be objectified by being unified and brought under a concept of judgement, Kant argued. But nothing in the content of the experience itself can accomplish this, so according to Kant's way of thinking the objectifying concept must be a priori, something which the mind itself brings to bear on the experiences from its own resources. Here Kant is describing what he takes to be a power of the human understanding, and not a power which some might have and some lack; and this is reflected in the sort of a priori synthesis which interests him. The content of any particular judgement is dependent on the individual's experiences; but its formal properties are bestowed by the understanding in its a priori aspect, which, being the same for all of us, guarantees the objectivity of the judgements we make in those respects to which the a priori categories apply. Therefore, according to Kant, we must see a cup, and the table it stands on, as external, three-dimensional and extended in time. For these are a priori conditions of our seeing them as objects at all. But the particular properties of the cup and the table are something which the individual experience is left to determine, and the objectivity or otherwise of the

judgements we make about these matters is something which, so far as Kant considers it at all, he leaves to science and common sense, operating within the framework of the categories. Thus Kant's account may be said to establish the necessary objectivity of the determinable features of the contents of the external world; and this is what results in the common belief that the claim that any judgement is objective amounts to the claim that what it says is necessarily true. But the Kantian account, as I understand it, has nothing to say about the determinate properties of objects, in respect of their objectivity or otherwise. For it has nothing to say about the objectivity of the determinate content of any judgement.

Kant's account therefore leaves the most compelling questions about objectivity unanswered. For judgements about the contingent properties of objects have just as much claim to be considered capable of being objective as judgements about the necessary properties of objects; but it is not easy to see how Kant would set about explaining this fact, because it is not easy to see how his usage of 'objective' could be made to stretch that far. Now since it is obvious that there are contingent judgements about objects which have an excellent claim to be objective — it is not just a commonly-held opinion that grass is green, for instance — ; and since it is also obviously true that disputes arise over what the contingent properties of objects are, we need to explain how the objectivity of the contingent content of judgements can be established. When two people disagree about the colour of something, for instance, and all the usual sources of error in such judgements have been discounted, we need to be able to explain which of their judgements (if either) is objective. (For until we have a method of determining that, we cannot in fact investigate and pin down 'the usual sources of error'.) From this too arises the problem of how we are to allow room, in any explanation we give, for the subjectivity in this respect of some of our judgements but not others.

My inability to find an answer to these questions in Kant's work was what first led me to adopt the line of enquiry I have taken. Not being a Kantian scholar, I may

have failed to detect at least the elements of an answer to these questions in Kant's work. But I do not think that is the case. For his reliance on the belief that reference to our common form of sensibility provides the general key to all epistemological problems makes it inevitable that the Critical Philosophy should have little to say about the contingent content of our judgements, or should consider the possibility of objectivity at that level of judgement at all. And this, I think, is as true of Kant's approach to the question of the objectivity of moral judgements as it is to his approach to empirical objectivity. Hence my extension of the same questions to the moral realm.

BIBLIOGRAPHY

Aristotle	Nichomachean Ethics (tr. J.A.K.Thomson, Penguin 1955)
Atkinson R.F.	Hume on 'Is' and 'Ought': a Reply to A.C. MacIntyre (Philosophical Review 1961)
Austin J.L.	Sense and Sensibilia (ed. G.J.Warnock, OUP 1962)
Ayer A.J.	Can there be a Private Language? (PAS 1954)
Ayer A.J.	The Problem of Knowledge (Penguin 1956)
Ayer A.J.	Logical Positivism (Glencoe 1959)
Ayer A.J.	Philosophical Essays (Macmillan 1963)
Bambrough R.	Moral Scepticism and Moral Knowledge (Routledge 1979)
Beardsmore R.W.	Moral Reasoning (Routledge 1969)
Beardsmore R.W.	Art and Morality (Macmillan 1971)
Berlin I. (ed)	The Age of Enlightenment (Mentor 1956)
Boyne C.W.	Vagueness and Colour Predicates (Mind 1972)
Brandt R.B.	Review of 'Freedom and Reason' (Journal of Philosophy 1964)
Bull N.J.	Moral Education (Routledge 1973)
Cassirer E.	The Philosophy of the Enlightenment (Princeton U.P. 1951)
Child D.	Psychology and the Teacher (Holt 1973)
Copi I.M.	Introduction to Logic (4th ed. Macmillan 1972)
Dummett M.	Wittgenstein's Philosophy of Mathematics (Philosophical Review 1959)
Edel M.& A.	Anthropology and Ethics (Case W.R.U. 1968)
Emmet E.R.	Learning to Philosophize (Penguin 1968)
Flew A.G.N.	On the Interpretation of Hume (Philosophy 1963)
Foot P.	Virtues and Vices (Blackwell 1981)
Foot P. (ed)	Theories of Ethics (OUP 1967)
Geach P.T.	Mental Acts (Routledge 1957)
Hacker P.M.S.	Insight and Illusion (Oxford 1972)
Hamlyn D.W.	Seeing Things As They Are (an Inaugural Lecture delivered at Birkbeck College 1965)
Hamlyn D.W.	Objectivity (in 'Education and the Development of Reason' ed. Dearden, Hirst and Peters, Routledge 1972)

Hare R.M. The Language of Morals (OUP 1952)

Hare R.M. Freedom and Reason (OUP 1963)

Hare R.M. Descriptivism (Proceedings of the British Academy 1963)

Hare R.M. Moral Thinking (OUP 1981)

Hudson W.D. Ethical Intuitionism (Macmillan 1967)

Hudson W.D. Modern Moral Philosophy (Macmillan 1970)

Hudson W.D. (ed) New Studies in Ethics Vols.I & II (Macmillan 1974)

Hume D. A Treatise of Human Nature (in Hume's Works, ed. Green & Grose: Scientia Verlag Aalen 1964)

Huxley A. The Doors of Perception (Penguin 1959)

Jackson R. Locke's Distinction between Primary and Secondary Qualities (Mind 1929)

Kant I. Critique of Pure Reason (tr. N.Kemp Smith, London 1929, Macmillan 1963)

Kant I. Groundwork of the Metaphysic of Morals (tr. H.J.Paton, London 1948)

Körner S. Kant (Penguin 1955)

Kuhn T.S. The Structure of Scientific Revolutions (2nd edition, Chicago 1970)

Locke J. An Essay Concerning Human Understanding (ed. A.S.Pringle-Pattison, OUP 1924, Harvester 1978)

MacIntyre A.C. Hume on 'Is' and 'Ought' (Philosophical Review 1959)

Mackie J.L. Problems from Locke (Oxford 1976)

MacKinnon D.M. Ethical Intuitionism (in 'Contemporary British Philosophy' (3rd series, ed. H.D. Lewis) Unwin 1956)

Malcolm N. Dreaming (Routledge 1959)

Malcolm N. Knowledge and Certainty (Englewood Cliffs 1963)

Mead M. Sex and Temperament in Three Primitive Societies (Routledge 1977)

Menen A. The Prevalence of Witches (Chatto & Windus 1970)

Moore G.E. Principia Ethica (Cambridge 1903)

Moore G.E. Ethics (OUP 1912)

Murdoch I.	Sartre (Bowes & Bowes 1953)
Nowell-Smith P.	Ethics (Penguin 1954)
O'Connor D.J.	John Locke (Penguin 1952)
Odegard D.	Two Arguments Against Moral Certainty (Mind 1981)
Passmore J.	A Hundred Years of Philosophy (Duckworth 1957)
Pears D.	Wittgenstein (Fontana 1971)
Perry T.D.	Moral Reasoning and Truth (Oxford 1976)
Peters R.S.	Ethics and Education (Unwin 1970)
Phillips D.Z. & Mounce H.O.	On Morality's Having a Point (Philosophy 1965)
Piaget J. & Szeminska A.	The Child's Conception of Number (Routledge 1952)
Pitcher G.	The Philosophy of Wittgenstein (Prentice-Hall 1964)
Popper K.	Conjectures and Refutations (Routledge 1963)
Prichard H.A.	Moral Obligation (Oxford 1949)
Putnam H.	Meaning and the Moral Sciences (Routledge 1978)
Rashdall H.	The Theory of Good and Evil (London 1924)
Russell B.	The Problems of Philosophy (OUP 1912)
Sartre J-P.	Existentialism and Humanism (tr. Mairet, Methuen 1948)
Smart J.J.C.	Philosophy and Scientific Realism (Routledge 1963)
Snare F.E.	The Diversity of Morals (Mind 1980)
Speake J. (ed)	A Dictionary of Philosophy (Pan 1979)
Toulmin S.	The Place of Reason in Ethics (CUP 1950)
Urmson J.O.	On Grading (Mind 1950)
Warnock G.J.	Contemporary Moral Philosophy (London 1967)
Warnock M.	Ethics Since 1900 (OUP 1960)
Westermarck E.	Ethical Relativity (Littlefield Adams 1960)
Wittgenstein L.	The Blue and Brown Books (Blackwell 1958)
Wittgenstein L.	Philosophical Investigations (tr. G.E.M. Anscombe, Blackwell 1958)
Wittgenstein L.	Lecture on Ethics (Philosophical Review 1965)
Wittgenstein L.	Zettel (tr. G.E.M.Anscombe, ed. G.E.M. Anscombe & G.H. von Wright, Blackwell 1967)
Wittgenstein L.	On Certainty (tr. D.Paul & G.E.M.Anscombe, ed. G.E.M.Anscombe & G.H. von Wright, Blackwell 1969)

Williams B. Morality: An Introduction to Ethics (CUP 1976)

Williams B. Philosophical Papers 1973-80 (CUP 1981)