1

What is the Moral Problem?

1.1 NORMATIVE ETHICS VS. META-ETHICS

It is a common fact of everyday life that we appraise each others' behaviour and attitudes from the moral point of view. We say, for example, that we did the *wrong* thing when we refused to give to famine relief, though perhaps we did the *right* thing when we handed in the wallet we found on the street; that we would be *better* people if we displayed a greater sensitivity to the feelings of others, though perhaps *worse* if in doing so we lost the special concern we have for our family and friends. Most of us take appraisal of this sort pretty much for granted. To the extent that we worry, we simply worry about getting it right.

Philosophers are ordinary folk. They too are concerned to get the answers to moral questions right. Indeed they have been so concerned that their attempts have found their way all the way into the philosophy classroom, as a casual glance at the philosophy offerings in almost any university will reveal. Subjects in normative ethics, subjects with names like 'Practical Ethics', 'Applied Ethics', 'Contemporary Moral Issues', 'Crime and Punishment', 'Ethics and the Environment', 'The Good Things of Life' and so on, are now taught nearly everywhere. Of course, philosophers tend to give more technical and systematic answers to normative ethical questions than ordinary folk. I doubt that anyone seeking advice from a friend has been taken through the standard utilitarian, deontological, or virtue theory line – unless, of course, the friend is a professional philosopher or a philosophy student. But this is only because philosophers tend to construct moral theories

to answer moral questions, whereas ordinary folk are happy to give and receive one-off answers. The important point, however, is that philosophers' theories do not generate answers that are different in kind to the answers ordinary folk give to moral questions. They are *merely* more technical and more systematic.

Despite their interest in normative ethics, however, philosophers have not tended to think that these sorts of questions are of the first importance in moral philosophy (though contrast Singer, 1973). Rather they have thought that we should do normative ethics only after we have given satisfactory answers to certain questions in meta-ethics. But why?

In meta-ethics we are concerned not with questions which are the province of normative ethics like 'Should I give to famine relief?', or 'Should I return the wallet I found in the street?', but rather with questions *about* questions like these. What does the 'should' in such questions mean? Does it signal that these questions are about some matter of fact? If so, then how do we justify giving one answer rather than another? In other words, what sort of fact is a moral fact? In what sense is moral argument simply a species of rational argument? And if the 'should' does not signal that moral questions are about a matter of fact then, again, how do we justify giving one answer rather than another to such questions? In other words, what is moral argument about? What is its point or function? What is the standard against which a good moral argument is to be measured?

As perhaps the final questions make clear, philosophers have surely been right to give meta-ethical questions a certain priority over questions in normative ethics. If moral argument is not simply a species of rational argument, then that calls into question the very role of moral argument in everyday life. For in everyday life such argument plays an integral role in decisions about how we are to distribute benefits and burdens. If these distributions of benefits and burdens cannot be *rationally* justified, then we must ask ourselves by what means, if any, they can be justified at all.

In terms of this distinction between normative ethics and meta-ethics I should now say that this book is unashamedly devoted to answering certain questions in meta-ethics. I defend two main claims. The first is that questions like 'Should I give to famine relief?', and 'Should I return the wallet I found in the

street?', are questions about a matter of fact, and that moral argument is therefore simply a species of rational argument, argument whose aim is to discover the truth. This first claim is argued for in an analytic spirit, by way of an analysis of what the 'should' in such questions means, an analysis that in turn depends on my second claim, which I will mention shortly. By way of anticipation, one striking feature of the analysis is that it not only legitimates the interest that philosophers take in normative ethics, but also makes taking such an interest crucial for the final resolution of meta-ethical questions.

1.2 META-ETHICS TODAY

Before saying anything else, however, it seems to me best to begin by acknowledging, and attempting to diagnose, the difficulties involved in giving any convincing answers at all to meta-ethical questions. For if one thing becomes clear by reading what philosophers writing in meta-ethics today have to say, it is surely that enormous gulfs exist between them, gulfs so wide that we must wonder whether they are talking about a common subject matter. Here is a sample.

We are told that engaging in moral practice presupposes that there exist moral facts, and that this presupposition is an error or mistake akin to the error of presupposition made by someone who engages in a religious practice when there is in fact no God (Mackie, 1977). And we are told that moral commitment involves no such error of presupposition; that moral talk happens inside a perfectly kosher practice (Brink, 1984; Blackburn, 1985a; McDowell, 1985; Nagel, 1986).

We are told that moral facts exist, and that these facts are ordinary facts, not different in kind from those that are the subject matter of science (Harman, 1986; Railton, 1986; Jackson, 1992). And we are told that moral facts exist, and that these facts are *sui generis* (Moore, 1903; McDowell, 1979; Sturgeon, 1985; Boyd, 1988; Brink, 1989; Dancy, 1993).

We are told that moral facts exist and are part of the causal explanatory network (Railton, 1986; Boyd, 1988; Sturgeon, 1985; Brink, 1989). And we are told not just that moral facts play no role

in the causal explanatory network, but that there are no moral facts at all (Ayer, 1936; Hare, 1952; Mackie, 1977; Blackburn, 1984; Gibbard, 1990).

We are told that there is an internal or necessary connection between moral judgement and the will (Hare, 1952; Nagel, 1970; Blackburn, 1984; McDowell, 1985; Korsgaard, 1986). And we are told that there is no such connection, that the connection between moral judgement and the will is altogether external and contingent (Frankena, 1958; Foot, 1972; Scanlon, 1982; Railton, 1986; Brink, 1986).

We are told that moral requirements are requirements of reason (Kant, 1786; Nagel, 1970; Darwall, 1983; Korsgaard, 1986). And we are told that it is not necessarily irrational to act immorally, that moral evaluation is different in kind from the evaluation of people as rational or irrational (Hume, 1888; Foot, 1972; Scanlon, 1982; Blackburn, 1984; Williams, 1985; Railton, 1986; Harman, 1985; Lewis, 1989).

We are told that morality is objective, that there is a single 'true' morality (Kant, 1786; Nagel, 1970; Darwall, 1983; Korsgaard, 1986; Brink, 1989). And we are told that morality is not objective, that there is not a single true morality (Ayer, 1936; Mackie, 1977; Harman, 1975, 1985; Williams, 1985).

Nor should it be thought that this account of the deep disagreements that exist is misleading; that though there are disagreements, there are certain dominant views. The situation is quite otherwise. There are no dominant views. In their recent comprehensive review of a century of meta-ethics, Stephen Darwall, Alan Gibbard and Peter Railton remark that the 'scene is remarkably rich and diverse' (1992: 124). But even to the casual observer, this is surely an understatement. The scene is so diverse that we must wonder at the assumption that these theorists are all talking about the same thing.

1.3 THE MORAL PROBLEM

Why do meta-ethical questions engender so much disagreement? In my view, the reason can be traced to two of the more distinctive

features of morality, features that are manifest in ordinary moral practice as it is engaged in by ordinary folk. The philosopher's task is to make sense of a practice having these features. Surprisingly, however, these features pull against each other, so threatening to make the very idea of morality altogether incoherent (Smith, 1991).

To begin, as we have already seen, it is a distinctive feature of engaging in moral practice that the participants are concerned to get the answers to moral questions *right*. And this concern itself seems to force certain meta-ethical conclusions. Such concern presupposes, for example, that there are correct answers to moral questions to be had. And the natural interpretation of that presupposition is that there exists a domain of moral facts; facts about which we can form beliefs and about which we may be mistaken.

Moreover, the way in which we conduct ourselves in living the moral life seems to presuppose that these facts are in principle available to all; that no one in particular is better placed to discover them than anyone else. That we have something like this conception of moral facts seems to explain our preoccupation with moral conversation and moral argument on the one hand, and novels and films in which the different reactions people have to moral questions are explored on the other. For what seems to give these preoccupations their point and poignancy is the simple idea that we are all in the same boat. A moral dilemma or issue faced by one person is a dilemma or issue that might be faced by another. Absent some relevant difference in their circumstances, what counts as an adequate response in one person's case must count as an adequate response in another's.

To put the point another way, we seem to think that the only relevant determinant of the rightness of an act is the circumstances in which the action takes place. If agents in the same circumstances act in the same way then either they both act rightly or they both act wrongly. Given that this is so we have the potential to learn something from each others' responses to the moral dilemmas and issues that we face. A careful mustering and assessment of the reasons for and against our particular moral opinions about such dilemmas and issues is therefore the best way to discover what the moral facts really are. If we are open-minded and

thinking clearly then such an argument should result in a convergence in moral opinion, a convergence upon the truth. Individual reflection may of course serve the same purpose, but only when it simulates a real moral argument; for only then can we be certain that we are giving each side of the argument due consideration; that we aren't simply congratulating ourselves on our own blind prejudices and preconceptions.

We may summarize this first feature of morality in the following terms: we seem to think moral questions have correct answers; that the correct answers are made correct by objective moral facts; that moral facts are wholly determined by circumstances; and that, by engaging in moral conversation and argument, we can discover what these objective moral facts determined by the circumstances are. The term 'objective' here simply signifies the possibility of a convergence in moral views of the kind just mentioned. Let's call this the 'objectivity of moral judgement'.

A second and rather different feature of morality concerns the practical implications of moral judgement. Suppose we are sitting together one Sunday afternoon. World Vision is out collecting money for famine relief, so we are waiting to hear a knock on the door. I am wondering whether I should give to this particular appeal. We debate the pros and cons of contributing and, let's suppose, after some discussion, you convince me that I should contribute. There is a knock on the door. What would you expect? I take it that you would expect me to answer the door and give the collector my donation. But suppose I say instead 'But wait! I know I *should* give to famine relief. But what I haven't been convinced of is that I *have any reason* to do so!' And let's suppose that I therefore refuse to donate. What would your reaction be?

It seems to me that your reaction would be one of extreme puzzlement. The conversation we had was about whether or not I should give to famine relief. But this just seems equivalent to a conversation about whether or not I have a reason to give to famine relief. Given that I claim to have been convinced by that conversation, and given that reasons have motivational implications, my refusal will therefore quite rightly occasion serious puzzlement. Perhaps I will be able to explain myself: perhaps I think that there is a better reason to do something else; or per-

haps I am suffering from weakness of will or some other such psychological failure and am claiming only that I haven't been convinced that I have been given any ulterior motive, a motive that will move me in the face of my own weakness. But absent some such explanation, the puzzlement will be such as to cast serious doubt on the sincerity of my claim to have been convinced that it is right to give to famine relief at all.

We can summarize this second feature of morality as follows: moral judgements seem to be, or imply, opinions about the reasons we have for behaving in certain ways, and, other things being equal, having such opinions is a matter of finding ourselves with a corresponding motivation to act. Let's call this the 'practicality of moral judgement'.

These two distinctive features of morality – the objectivity and the practicality of moral judgement – are widely thought to have both metaphysical and psychological implications. However, and unfortunately, these implications are the exact opposite of each other. In order to see why this is thought to be so, we need to pause for a moment to reflect more generally on the nature of human psychology.

According to the standard picture of human psychology – a picture we owe to Hume (1888) – there are two main kinds of psychological state. On the one hand there are beliefs, states that purport to represent the way the world is. Since our beliefs purport to represent the world, they are assessable in terms of truth and falsehood, depending on whether or not they succeed in representing the world to be the way it really is. And on the other hand there are desires, states that represent how the world is to be. Desires are unlike beliefs in that they do not even purport to represent the way the world is. They are therefore not assessable in terms of truth and falsehood. Hume concludes that belief and desire are therefore distinct existences: that is, that we can always pull belief and desire apart, at least modally. For any belief and desire pair that we imagine, we can always imagine someone having the desire but lacking the belief, and vice versa. If this were not so, if we had to imagine a particular belief bringing a particular desire with it, then desires would – contrary to fact – be assessable in terms of truth and falsehood, at least derivatively; for

we could count the desire as true whenever the belief with which it was necessarily connected counted as true, false whenever the belief with which it was necessarily connected counted as false.

Indeed, according to the standard picture of human psychology that we get from Hume, not only are desires not assessable in terms of truth and falsehood, they are not subject to any sort of rational criticism at all. The fact that we have a certain desire is, with a proviso to be mentioned presently, simply a fact about ourselves to be acknowledged. It may be unfortunate that we have certain combinations of desires – perhaps our desires cannot all be satisfied together – but in themselves our desires are all on a par, rationally neutral. This is an important further claim about desires. For it suggests that though we may make discoveries about the world, and though these discoveries may rightly affect our beliefs, such discoveries should, again with one proviso to be mentioned presently, have no rational impact upon our desires. They may, of course, have some *non*-rational impact. Seeing a spider I may be overcome with a morbid fear and desire never to be near one. However this is not a change in my desires mandated by reason. It is a *non*-rational change in my desires.

Now for the proviso. Suppose, contrary to the example I just gave, that I acquire the desire never to be near a spider because I come to believe, falsely, that spiders give off an unpleasant odour. Then we would certainly ordinarily say that I have an 'irrational' desire. However the reason we would say this clearly doesn't go against the spirit of what has been said so far. For my desire never to be near a spider is *based on* a further desire and belief: my desire not to smell that unpleasant odour and my belief that that odour is given off by spiders. Since I can be rationally criticized for having the belief, as it is false, I can be rationally criticized for having the desire it helps to produce. The proviso is thus fairly minor: desires are subject to rational criticism, but only insofar as they are based on beliefs that are subject to rational criticism. Desires that are not related in some such way to beliefs that can be rationally criticized are not subject to rational criticism at all.

According to the standard picture, then, there are two kinds of psychological state – beliefs and desires – utterly distinct and different from each other. The standard picture of human psy-

chology is important because it provides us with a model for explaining human action. Crudely, our beliefs tell us how the world is, and thus how it has to be changed, so as to make it the way our desires tell us it is to be. An action is thus the product of these two distinct existences: a desire representing the way the world is to be and a belief telling us how the world has to be changed so as to make it that way.

Let's now return to consider the two features of moral judgement we discussed earlier. Consider first the objectivity of such judgement: the idea that moral questions have correct answers, that the correct answers are made correct by objective moral facts, that moral facts are determined by circumstances, and that, by engaging in moral argument, we can discover what these objective moral facts are. The metaphysical and psychological implications of this may now be summarized as follows.

Metaphysically, the implication is moral realism: the view that, amongst the various facts there are in the world, there aren't just facts about (say) the consequences of our actions on the well-being of our families and friends, there are also distinctively moral facts: facts about the rightness and wrongness of our actions having these consequences. And, psychologically, the implication is thus cognitivism: the view that when we make a moral judgement we thereby express our beliefs about the way these moral facts are. In forming moral opinions we acquire beliefs, representations of the way the world is morally.

Given the standard picture of human psychology, there is a further psychological implication. For whether or not people who have a certain moral belief desire to act accordingly must now be seen as a further and entirely separate question. They may happen to have a corresponding desire, they may not. However, either way, they cannot be rationally criticized. Having or failing to have a corresponding desire is simply a further fact about a person's psychology. On this view, believing that, say, I should give to famine relief does not require that, other things being equal, I have a reason to give to famine relief.

But now consider the second feature, the practicality of moral judgement. We saw earlier that to have a moral opinion simply is, contrary to what has just been said, to find ourselves with a corre-

sponding motivation to act. If we think it right to give to famine relief then, other things being equal, we must be motivated to give to famine relief. The practicality of moral judgement thus seems to have a psychological and a metaphysical implication of its own.

Psychologically, since making a moral judgement requires our having a certain desire, and no recognition of a fact about the world could rationally compel us to have one desire rather than another, our judgement must really simply be an expression of that desire, or perhaps a complicated disposition to have that desire. This is non-cognitivism. And this psychological implication has a metaphysical counterpart. For it seems to follow that, contrary to initial appearance, when we judge it right to give to famine relief we are not responding to any moral fact: the rightness of giving to famine relief. Indeed, moral facts are an idle postulate. In judging it right to give to famine relief we are really simply expressing our desire, or disposition to desire, that people give to famine relief. It is as if we were yelling 'Hooray for giving to famine relief!' No mention of a moral fact there. Indeed, no factual claim at all. This is irrealism: the denial of the claim that there are any moral facts. But importantly it is irrealism in the form of expressivism: the view that in making moral judgements we do not even purport to make claims about how things are morally, but rather simply give expression to our non-cognitive states.

On this view, when I claim that I should give to famine relief, it does indeed follow that I have a reason to give to famine relief, at least other things being equal. That follows from the fact that I give expression to my desire, or disposition to desire. And when I claim that I should give to famine relief, it follows that I commit myself to judging that it is right for anyone in circumstances like mine to give to famine relief as well. That follows from the scope of my desire. But I need not think that these other people have any reason to give to famine relief. And, moreover, I should not think that the claim that people in circumstances like mine should give to famine relief is itself objective. For there is no reason to suppose that others could be brought to agree with me by means of a rational argument. A failure to elicit agreement from others would not necessarily be evidence of error or mistake

or confusion on someone's behalf. It would simply be evidence of the fact that they have different desires, or dispositions to have desires, from those that I have.

We are now in a position to see why meta-ethical questions engender so much disagreement. The task of the philosopher in meta-ethics is to make sense of ordinary moral practice. But the problem is that ordinary moral practice suggests that moral judgements have two features that pull in quite opposite directions from each other. The objectivity of moral judgement suggests that there are moral facts, wholly determined by circumstances, and that our moral judgements express our beliefs about what these facts are. This enables us to make good sense of moral argument, and the like, but it leaves it entirely mysterious how or why having a moral view is supposed to have special links with what we are motivated to do. And the practicality of moral judgement suggests just the opposite, that our moral judgements express our desires. While this enables us to make good sense of the link between having a moral view and being motivated, it leaves it entirely mysterious what a moral argument is supposed to be an argument about; the sense in which morality is supposed to be objective.

The idea of morality thus looks like it may well be incoherent, for what is required to make sense of a moral judgement is a strange sort of fact about the universe: a fact whose recognition necessarily impacts upon our desires (Mackie, 1977). But the standard picture of human psychology tells us that there are no such facts. Nothing could be everything a moral judgement purports to be. Perhaps we should therefore all be irrealists, and deny the existence of moral facts. But we should combine our irrealism with cognitivism: the view that our moral judgements purport to make claims about moral facts. In other words, we should all be moral nihilists, and simply acknowledge that moral practice is founded on a massive error of presupposition. Or so it may now seem.

This is what I call 'the moral problem' (Smith, 1989; 1994a). As I see it, the moral problem is in fact the central organizing problem in contemporary meta-ethics. It explains the massive disagreement that exists among philosophers about meta-ethical issues.

This problem can be stated succinctly in the form of three apparently inconsistent propositions (see also McNaughton, 1988: 23).

1 Moral judgements of the form 'It is right that I ϕ' express a subject's beliefs about an objective matter of fact, a fact about what it is right for her to do.

2 If someone judges that it is right that she ϕs then, *ceteris paribus*, she is motivated to ϕ.

3 An agent is motivated to act in a certain way just in case she has an appropriate desire and a means-end belief, where belief and desire are, in Hume's terms, distinct existences.

The apparent inconsistency can be brought out as follows: from (1), the state expressed by a moral judgement is a belief, which, from (2), is necessarily connected in some way with motivation; that is, from (3), with having a desire. So (1), (2) and (3) together entail that there is some sort of necessary connection between distinct existences: moral belief and desire. But (3) tells us that there is no such connection. Believing some state of the world obtains is one thing, what I desire to do given that belief is quite another.

In light of the moral problem it is easy to understand not just why meta-ethical questions engender so much disagreement, but also why these disagreements take just the form that they do. For those who offer answers to meta-ethical questions must now make a choice. If they are to remain faithful to the standard Humean picture of human psychology, then they must decide which of the two features of moral judgement – objectivity and practicality – to reject, and they then have to defend their choice. Thus it should come as no surprise that we find the expressivists rejecting (1), the claim that moral judgements express beliefs (Ayer, 1936; Hare, 1952; Blackburn, 1984, 1986, 1987; Gibbard, 1990), the externalists rejecting (2), the claim that there is a necessary connection of sorts between moral judgement and motivation (Frankena, 1958; Foot, 1972; Scanlon, 1982; Railton, 1986; Brink, 1986, 1989), and the anti-Humean theorists of motivation breaking ranks altogether and rejecting (3), the claim that motivation

is to be explained in terms of a desire and means-end belief, where belief and desire are, in Hume's terms, distinct existences (Nagel, 1970; McDowell, 1978; Platts, 1981; McNaughton, 1988; Dancy, 1993). Moreover, it should come as no surprise that each of these theorists – the expressivists, the externalists, and the anti-Humean theorists of motivation – argue that the proposition that they reject, as opposed to the propositions the others reject, is a mere philosophical fantasy; something to be explained away. But, as is now I hope plain, no matter which proposition these philosophers choose to reject, they are bound to end up denying something that seems more certain than the theories they themselves go on to offer (Smith, 1989). Moral nihilism quite rightly looms (Smith, 1993a).[1]

1.4 TOWARDS A SOLUTION TO THE MORAL PROBLEM

The stage is thus set for this book. For in my view we are not forced to choose between accommodating the objectivity and the practicality of moral judgement in giving answers to meta-ethical questions. The correct solution to the moral problem allows us to accommodate both. Moreover, it does so without requiring us to give up the standard account of the explanation of action in terms of belief and desire and without plunging us into moral nihilism either.

I said earlier that I will defend two claims. The first is that questions like 'Should I give to famine relief?' are questions about a matter of fact, facts that we can discover through rational argument, argument whose aim is to discover the truth. The second claim is that, contrary to the Humean view, moral facts are indeed facts about the reasons that we all share. As I see it, we can accommodate both the objectivity and the practicality of moral judgement – propositions (1) and (2) that make up the moral problem – by defending these two claims. The trick is to show how these two claims, and the second in particular, can be made to square with the standard Humean account of the explanation of human action: that is, with proposition (3) that makes up the

moral problem. How much of the whole psychological theory that we have inherited from Hume can we accept, how much must we reject? In essence that is, if you like, the major sub-plot of this book.

My argument will turn on a distinction between two kinds of reasons: motivating and normative. I will argue that we can and should accept Hume's claim that belief and desire are distinct existences, states with distinctive roles to play in the explanation of human action. For this is a true and illuminating claim about our motivating reasons: the reasons that motivate, and thus explain, our actions. However, I will argue, we cannot and should not accept Hume's claim that desires are themselves beyond rational criticism. For that amounts to a false claim about our normative reasons: the reasons that rationally justify our actions. It is therefore in the formulation of a distinctively anti-Humean theory of normative reasons that we will find room for the idea of moral facts as facts about the reasons we all share; it is here that we will find the solution to the moral problem. Or so I will argue.

1.5 SUMMARY AND PREVIEW

In this chapter I have explained what the difference is between normative ethical questions and meta-ethical questions. Moreover, I have defended the view, common among philosophers, that we should begin our study of ethics by focusing on meta-ethics, not normative ethics. For we cannot hope to do normative ethics without first knowing what the standards of correct argument in normative ethics are, and it is in meta-ethics that we discover these standards.

However, as we have seen, meta-ethical questions engender much disagreement. I have argued that this is because there is a central organizing problem in meta-ethics. Can we reconcile the objectivity and the practicality of moral judgement with the standard picture of human psychology that we get from Hume? This is what I call 'the moral problem'. My suggestion is to be that we can solve this problem by defending an analysis of what the 'should' means in moral questions, questions like 'Should I give to

famine relief?' According to this analysis, a positive answer to this particular question implies that there is a reason for anyone in circumstances like mine to give to famine relief, a reason to be discovered through rational argument.

However, before giving my own view I begin by discussing the more standard solutions to the moral problem. In chapter 2, the expressivists' challenge to the claim that our moral judgements express our beliefs about an objective matter of fact is considered and rejected. In chapter 3, I consider and reject the externalists' challenge to the claim that there is some sort of necessary connection between moral judgement and motivation. And in chapter 4, I consider and reject the anti-Humean challenge to the standard, Humean, account of the explanation of action in terms of belief and desire, Hume's account of motivating reasons.

By the end of chapter 4, then, the standard solutions to the moral problem will all have been shown to be *bad* solutions. The claims that these standard solutions choose to reject should not be rejected at all. The remaining two chapters are more constructive.

In chapter 5, I develop an account of reasons – normative reasons – that allows us to say why certain desires are rationally required, and I explain why this account of normative reasons is not inconsistent with the Humean account of motivating reasons defended in chapter 4. And in chapter 6, I show how we can use this account of normative reasons to analyse our concept of rightness, and I explain how this analysis of rightness in terms of normative reasons enables us to provide a solution to the moral problem.