

MORALITY AND THE AMORAL AGENT

by

GREGORY WILLIAM DURWARD

B.Math., University of Waterloo, 1969

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

in

THE FACULTY OF GRADUATE STUDIES
(Department of Philosophy)

We accept this thesis as conforming
to the required standard

THE UNIVERSITY OF BRITISH COLUMBIA

September, 1978

© Gregory William Durward, 1978

In presenting this thesis in partial fulfilment of the requirements for an advanced degree at the University of British Columbia, I agree that the Library shall make it freely available for reference and study.

I further agree that permission for extensive copying of this thesis for scholarly purposes may be granted by the Head of my Department or by his representatives. It is understood that copying or publication of this thesis for financial gain shall not be allowed without my written permission.

Department of Philosophy

The University of British Columbia
2075 Wesbrook Place
Vancouver, Canada
V6T 1W5

Date

April 9 / 1979

ABSTRACT

The thesis is an exploration of the possibility and rationality of amoral agency. A characterization of the moral agent and his mode of practical deliberation is developed, taking as central the objectivity or impersonal validity of moral judgments and principles, and the concern of morality with the welfare of persons in general. This provides the framework for a discussion of two main forms of amorality.

A person may qualify as an amoral agent either because he has little or no concern for other persons' well-being and recognizes no valid claim on him to thus concern himself, or because he operates with a radically subjectivist view of practical reasons and principles of action. As an interesting and widely discussed example of the former sort of amoralism, a number of forms of egoism are discussed and it is argued that, while ethical egoism is untenable, there are individualistic egoistic theories which avoid the most serious anti-egoist arguments.

The more interesting sort of amorality, from a philosophical point of view, is the subjectivist one and some effort is made to outline a subjectivist theory of practical reasoning. Subjectivist metamoral theories are considered and rejected as accounts of the ordinary moral consciousness. The possibility of treating such theories as revisionistic in nature is discussed and the conclusion is that, while they may underwrite an outlook which can plausibly be regarded as moral, the subjectivist moral agent can maintain

his position only with difficulty. This is because there is a natural drift from moral thinking to an objectivist posture and vice versa.

CONTENTS

| | |
|---|-----|
| INTRODUCTION | 1 |
| I MORALITY AND THE MORAL AGENT | 10 |
| 1 The moral agent introduced | 10 |
| 2 The objectivity of moral judgments | 15 |
| 3 The welfare of others | 22 |
| 4 The moral point of view | 26 |
| 5 Some other features of morality and moral agency | 30 |
| 6 The amoral agent | 39 |
| II EGOISM | 43 |
| 1 Egoism as a form of amorality | 43 |
| 2 Morality and self-interest | 44 |
| 3 The concept of self-interest | 48 |
| 4 Ethical egoism - an objectivist amorality | 50 |
| 5 Personal egoism - another objectivist amorality | 55 |
| 6 Kurt Baier | 59 |
| 7 David Gauthier | 64 |
| 8 Subjectivist egoism | 73 |
| 9 The rationality of egoism | 75 |
| III THE GROUNDS OF MORAL AGENCY | 78 |
| 1 The possibility of amorality - non-rational capacities | 78 |
| 2 The possibility of amorality - conceptual structures and rationality | 84 |
| 3 Does "Why be moral?" make sense? | 88 |
| 4 The Golden Rule | 94 |
| 5 Thomas Nagel | 103 |

| | | |
|----|---|-----|
| IV | OBJECTIVISM AND SUBJECTIVISM | 125 |
| 1 | Objectivism | 125 |
| 2 | Subjectivism | 133 |
| 3 | The error theory | 139 |
| 4 | Against objectivism | 146 |
| 5 | The problems in metamoral subjectivism | 154 |
| 6 | Revisionist theories | 160 |
| V | REASONS AND VALUES - A SUBJECTIVIST OUTLOOK | 167 |
| 1 | Facts, reasons, and the definition of "moral" | 167 |
| 2 | Reasons and explanation | 174 |
| 3 | Reasons and justification | 178 |
| 4 | Practical debate and vocabulary | 184 |
| 5 | What can be valued? | 187 |
| 6 | Valuing and commending | 191 |
| 7 | Facts and reasons again | 197 |
| VI | MORAL AND AMORAL SUBJECTIVISM | 203 |
| 1 | The individual moral subjectivist | 203 |
| 2 | Forms of moral subjectivism | 204 |
| 3 | Moral and amoral subjectivism contrasted | 212 |
| 4 | Existentialism | 214 |
| 5 | Jean-Paul Sartre | 216 |
| 6 | Nietzsche | 220 |
| | BIBLIOGRAPHY | 226 |

ACKNOWLEDGEMENTS

The responsibility for the world's being burdened with yet another PhD must be borne chiefly by D.G. Brown, my supervisor, and M. Lynne Durward, my spouse. Without their patient support this thesis would undoubtedly have met the fate it so often seemed, in my eyes, to deserve. I am especially indebted to Dr. Brown for his helpful comments on my work throughout the past three years. Stimulating discussions with Edwin Levy, Gary Wedeking, Alistair Browne and Elbridge Rand also provided valuable food for thought. Finally, my thanks to my diligent and cheerful typist, Bev Thompson.

INTRODUCTION

Everyone will agree, I think, that there are people who, for one reason or another, cannot be classed as moral agents. Persons whose mental development has been severely retarded or whose mental functions have been highly disordered or impaired through accident or disease are among the clearest cases. Such people are exempt from moral criticism; we neither praise them nor blame them for their actions. In extreme cases we are strongly inclined to think of them in roughly the way we think of vegetables or animals, at least as far as their responsibility for their actions (if any) is concerned. We may suppose that we have moral obligations to them, of course, but we do not think they have obligations of a similar or correlative sort.

It is not easy to account for our attitudes towards extreme cases of the kind just mentioned. In particular, why do we exempt them morally? Is it because they lack the concepts of right and wrong, of obligation and so on, because they are not rational beings, because they cannot form or maintain a stable and coherent view of their situation, or is it some combination of these and other things? Perhaps the best we can hope for by way of a general answer is to note that there are certain capacities which someone must have in order to qualify as a moral agent. While there is considerable room for debate on just what these capacities are and to what degree and in what combination they

must be present, we can expect that things like a reasonable amount of general intelligence, a reasonably reliable memory, a fair degree of rationality in means/end thinking, a fairly "normal" view of the nature of their environment, and so on, will be central. There may also be some constraints concerning their emotional makeup.

In any case, I think that the consideration of capacities holds some promise of enabling us to distinguish those entities which are non-moral from those which are moral, in one very broad sense of "moral." Roughly, what we can expect is a bifurcation of entities into inanimate objects, the animals and plants, and certain sorts of "defective" human beings on the one hand, and rational, relatively normal mature adults on the other. There are bound to be very difficult borderline cases involving not only questions of degree of capacity development, but also questions about whether it is actual or potential capacities which are to count. What about children, the senile, the temporarily disabled, and the mentally disturbed? I am not very much interested in these difficult questions here, however. All that is required is to note that there is a distinction which can, albeit with difficulty, be drawn between the moral and the non-moral in terms of those capacities, whatever they are, which are required for someone to form moral concepts, participate in moral debate, deliberate morally, appreciate moral reasons, and so on. To be non-moral is to exist outside of the moral by being *incapable* of participating in the moral realm. By contrast, to be amoral is to exist outside of the moral in some other way.

I think that there are people who meet all the prerequisites for moral thinking and yet who fail to qualify as moral agents in another sense of "moral." These people are not non-moral agents. They are relatively normal human beings but they do not operate with the normal moral perspective. There are features of the ordinary moral consciousness which are conspicuously absent from their practical thinking. What these features are will be investigated in due course. Such people, and I realize at this stage little has been done to properly identify them, I will call amoral agents. Unless the context indicates otherwise, I will be using the term "moral agent" in the sense in which it contrasts with "amoral agent."

There are many sorts of amoral agents, as we will see, but there is one thing which can be said to restrict the range of persons with which I will be concerned. I am not here interested in persons who are amoral only because their environment is such that there is no opportunity or no need for them to think morally. The prisoner, the hermit, the castaway and the solitary scientist may all lead lives largely devoid of moral concerns but it would be appropriate to say that they are still moral agents as long as they *would* begin to appreciate things morally if their circumstances were to change.

So far I have contrasted "moral" with "non-moral" and with "amoral." If we can say that these contrasts correspond to two senses of the term "moral," then "moral" has also a third sense in which it contrasts with "immoral." The judgment that some person, or some action of a person, is immoral is clearly an

evaluative one and can only be made sincerely by someone who operates within a certain conceptual framework. "Immoral" means roughly "contrary to the requirements of morality" and therefore can only be used sincerely (ignoring inverted comma uses) by someone who recognizes the possibility of there being requirements of a certain sort on actions. The judgment is properly a moral one and therefore it will be misleading if it is made by an amoral (or even non-moral) agent, unless this is known and taken into account. But not only might it be misleading for an amoral agent to use the term "immoral," it could also be misleading for a moral agent to use it to describe an amoralist or some action of his. Calling someone immoral suggests that he is open to certain sorts of considerations and certain modes of discussion and persuasion (moral ones). It is the amoralist's unsusceptibility to these considerations and his abstention from moral forms of thinking which make it misleading to say that he acted immorally. He acted, we should say *outside* of moral considerations and not so much *against* them.

Unfortunately I find it difficult to say much more at this point to help make clear just what sort of persons I am interested in and whom I call amoral agents. The attempt to do so will be the ongoing project in virtually all of what follows. I trust that my claim that there is an interesting category of persons who are not non-moral and yet who are amoral, will be vindicated in due course as the nature of amoral agency is discussed.

Before proceeding, something should perhaps be said by way

of providing a rationale for some of the structure of the thesis and of explaining why I deal with some of the topics I do. Let me attempt to do this by reconstructing briefly the way in which I came to see amorality as a subject deserving consideration. Hopefully this will help to provide the reader with something of a perspective on what follows.

I have long been attracted by the view that the very best, if not ultimately the only, reasons a person could have for doing something are reasons of self-interest. Thus, I began by asking if some form of egoistic theory in ethics might be found which allows such reasons a central place but which at the same time does justice to the facts of human motivation. It seemed clear enough that no straightforward form of *psychological* egoism, (according to which no one does or can ever act except out of self-interest) could be true without a gross perversion of the concept of self-interest. Yet I had hoped that under a sufficiently sophisticated version of the notion of a person's interest, some formulation of *ethical* egoism might prove plausible enough.

Like many people I held the view that there is a correct ethical theory. I supposed that certain sorts of things are "really" reasons for acting in certain ways and further that some reasons are "really" better than others. I believed that there are correct normative principles with the help of which it is in principle possible to determine which solutions to a practical problem are correct and which are incorrect, given, of course a full acquaintance with the facts of the case. Human beings, being identical in their essential nature, must all be subject

to the requirements of the same normative theory.

My attempts to develop a tenable form of egoism, however, led me into a number of difficulties. First, the notion of self-interest is very restrictive. How, for example, could one have a good reason to be honest or to keep promises under imaginable situations in which one could do better by lying or breaking faith? And is it really plausible to suppose that under such circumstances no one could possibly have a very good reason to tell the truth or keep a promise?

More importantly, however, the literature on egoism is full of "refutations" which purport to show that ethical egoism is beset with either practical or theoretical inconsistencies. Since the ethical egoist holds that everyone ought to pursue his own self-interest and since his securing advantages for himself will sometimes require that others be less than enthusiastic about pursuing their interests, the egoist must hold a theory which he cannot advocate. This is enough to show that there are grave problems in the way of holding ethical egoism as a moral theory. I will have a good deal more to say about egoism later but I simply want to indicate here how one might be driven to take seriously the possibility of adopting a practical theory (a theory about the nature of one's reason) which is not a moral theory.

The more I thought about the rationality of the egoist's position, the less peculiar the idea of a non-moral practical theory seemed. If the egoist could rationally remain outside morality, as it seemed he could, might there not be other forms

of amorality perhaps more radical and more interesting and no less rational?

Having come to the point of being prepared to take amorality seriously, it became clear that I needed to generate a characterization of morality, moral thinking and moral agency in order to explore systematically the ideas of amorality, amoral thinking and amoral agency. The results of this part of the investigation appear in Chapter I. There I argue that most of what seems right in the plethora of metamoral theories which have been advanced can be understood in terms of two basic features of moral thinking; an objectivist posture regarding questions of value and a concern for the welfare of other persons. The next chapter, on egoism, then constitutes an exploration of one form of amoralism; viz., one in which no (direct) concern for the welfare of other persons is present.

In spite of the fact that some forms of egoism are, I argue, logically and psychologically possible positions, they are nonetheless somewhat peculiar. The egoist's concerns are very narrowly delimited indeed, and this robs the egoistic orientation of much of its real interest. It would be much more instructive if forms of amorality could be found which allow a rather broader range of interests. But before proceeding to attempt this I found it useful to pause to consider some of the arguments which have been advanced in support of moral agency, and hence against amoral agency. Chapter III is a survey of some of these arguments and it simultaneously focusses the objectivity issue somewhat. It is in the adoption of a radically subjective approach to value

questions that I attempt to locate a more interesting kind of amorality.

Chapter IV is devoted to an exploration of the concepts of objectivity and subjectivity. There I argue that subjectivist metamoral theories are unacceptable as attempts to describe the phenomena of ordinary moral thinking and discourse. Some problems in objectivist approaches are discussed. Chapter V is a tentative account of an alternative to the objectivism which I claim is evident in, and central to, moral agency. Finally, in Chapter VI, I suggest that while a subjectivist outlook is not strictly entirely inconsistent with moral agency, the form of subjectivism which can be adopted by a moral agent can be maintained only with considerable effort and with the help of a certain self-conception. Since this self-conception differs from that of an amoralist, there is still a possibility of distinguishing the moral from the amoral agent even when we consider subjectivist versions of each. This chapter concludes with a brief consideration of J.P. Sartre's existentialist theory and Nietzsche's concept of *das Uebermensch* and their relevance to amorality as it is developed in the thesis.

I hope to show that it is a real and rational option for a mature and reflective person to stand outside of morality in certain ways. The most philosophically interesting form of amorality and hence the one which receives most attention is the one which involves an extremely individualistic and subjectivist outlook on the nature of practical reasoning. It is also to be hoped that the exploration of amoral agency will shed some useful

light on the concept of moral agency in roughly the way that attention to exceptions helps to clarify the rule. The possibility of amorality is not often taken very seriously by moral philosophers and I know of no systematic exploration of the idea. Too often it is simply assumed either that there *must* be good arguments against anyone who fails to appreciate things morally (if only we can find them) or worse, that there is no use arguing with such a person at all. In some cases these assumptions may be well justified but in general they are based on no real understanding of the amoral outlook. Perhaps what follows will serve as a beginning toward that understanding.

I

MORALITY AND THE MORAL AGENT

1 The moral agent introduced

Since I am concerned to explore the nature of a class of persons who operate, in some sense, outside of morality, the natural place to begin is with an attempt to characterize a moral agent. Unfortunately, this is not at all a simple matter, largely because the concept of morality is not one which can be analyzed in a very straightforward way. Many attempts have been made to get clear on just what is involved in thinking morally and to give some account to the meaning of moral terms, the nature of moral debate, the character of moral principles, and so on. Yet in spite of the amount of energy which has been expended on these and related problems, nothing like a clear consensus has even begun to emerge. I certainly do not wish to claim that the following constitutes a resolution of these difficult problems, but then I am not sure that the quest for a definitive analysis of moral language is one which holds much promise of success anyway. As is often the case in philosophy, the real value of a line of enquiry lies, not in the final answer, but in the process of exploring the issues which arise along the way. Even if we seem to end where we started, we can hope that something of interest has been glimpsed en route. Round trip fares are not even usually a waste of money.

Let me begin the attempt to characterize the moral agent by simply producing a composite sketch whose elements are drawn from the testimony of philosophers who have given the problem their careful consideration. If the end product is a recognizable facsimile of a type of person we cannot but call a moral agent, some progress will have been made, especially if it enables us to construct a portrait of his counterpart, the amoralist.

Throughout this chapter there will be something of a problem involved in keeping clear the distinction between the moral agent (the not-moral agent) and the morally good agent (the not-immoral agent). This is because the distinction between the morally good person and the immoral (or morally inferior) person is largely a matter of how consistently and how successfully a person operates as a moral agent - of how seriously and conscientiously he takes his involvement in the moral realm. As an analogy, consider a man playing a game of golf. He may play following the rules to the letter, conscientiously noting his strokes on each hole, playing the ball from where it lies each time, and so on. But another man may cheat, breaking or bending the rules, either to win or to avoid certain frustrations which detract from his enjoyment of the game. Insofar as this second man is playing golf, however, he will typically react in one or more of certain predictable ways if confronted with his transgressions. He may deny that he has committed them; he may admit them and attempt to justify them in some way, e.g. by an appeal to some allegedly higher principle or purpose ("After all we're here to enjoy ourselves") or he may go so far as to advocate that the rules of the game be changed to conform with

his practice. Such a man is playing golf and even though he is not always, or carefully, following the rules, he recognizes the validity of those rules.

Contrast these cases with one in which a man goes from hole to hole on a golf course, sometimes hitting the ball in the "correct" manner, sometimes throwing it and sometimes kicking it. Occasionally he skips a hole or plays one twice, and so on. Confronted with the charge of cheating or unsportsmanlike conduct, he is quite unperturbed. His only answer is that he isn't interested in all those rules and constraints because he isn't playing golf. This man is not cheating at golf - he isn't playing golf. The first man we considered is like the morally good man and the second like the immoral man. Both are, in some sense, playing the game of morality. The immoral man will typically deny, justify, excuse or defend his actions when confronted with the charge of moral wrongdoing. He recognizes the validity of morality and accepts the right of others to require of him either an excuse, a justification or an admission of guilt and the acceptance of their right to reprimand or punish him. The third man, the one who is not playing the game of golf, is like the amoral agent. Properly speaking he isn't cheating or acting immorally, unless participation in the game is not voluntary. (Actually that is the point at which the analogy is weakest because most people, I think, do not suppose that the requirements of morality are escapable in anything like the way the rules of golf are. If you *can* play, you *must* play, like it or not, and your only choice is to play well or poorly.)

In any case, in describing the rules of golf, one can largely

ignore the fact that some people cheat, and similarly in describing the nature of morality or moral agency, one may overlook the fact that people often act immorally. This is because the immoral man admits, in some sense, the validity of moral requirements, however inconsistent and lax he may be in living those claims, and however perverse (within limits of course) his interpretation of them. He does not become amoral, for example, by convincing himself that his own case is a special one exempt from the usual rules: on the contrary, in his very effort to make of his own case an exception, he is engaged in the moral enterprise.

Let me offer a thumbnail sketch of the moral agent. (The features of the sketch will be elucidated and expanded in subsequent sections.)

The moral agent is a participating member of a social unit and he sees himself as such. Confronted with a practical problem in which the interests or welfare of other persons in his social group are involved, he considers the effects of his actions on these persons, not because it is to his benefit to do so but because he recognizes a requirement that he do so (and ideally because he cares about them). He recognizes duties and obligations which function as valid claims on him and which restrict his freedom to do as he pleases. In situations in which the welfare of others is involved, he asks himself what he ought to do, what actions would be right or wrong to do, and so on, and in doing so he supposes that, in some fairly strong sense, there are correct or incorrect ways of resolving the practical problems he faces. He must decide what to do but his decision can be mistaken since there is something

like a fact of the matter about the question of what he ought to do.

In deliberating, he takes up the moral point of view from which he sees himself as just one person among many and from which his own first-order interests appear pretty much on a par with the first-order interests of others. From this point of view he can be assured of arriving at a resolution which is acceptable to anyone and everyone since it is most likely to represent a solution unbiased by personal tastes and desires.

He uses moral concepts and participates in moral debate with others, requiring of others and providing as best he can, justifications for various lines of conduct. His reasons and justifications are, he supposes, assessable in principle by interpersonally valid standards. His moral judgments and those of others, he treats as if they have a determinate truth value. Moral determinations have for him, and ought to have for all, he thinks, an overridingness and authority in practical matters. Whatever the outcome of a situation, its value is unaffected by who is considering the matter. There are questions of value which cannot be assimilated to matters of preference or taste, however deep and pervasive.

Now I suppose that there will hardly be a single philosopher who could not find something to quarrel with in this characterization, and there are those who would reject a good deal of it outright. Still, a start must be made somewhere and it has been my experience that, radical as it might seem, something very like the above will be accepted by many and perhaps most people as approximating the way they view themselves. I do not wish to minimize the importance,

however, of the fact that there are different - indeed radically different - ways of characterizing the moral agent. But it would be a mistake, I think, to suppose that what follows will have little or no interest to someone who does not subscribe to the analysis I adopt. The issues I will be dealing with will, in large measure, be issues for anyone interested in the nature of practical reasoning, even if the terms in which they are discussed here seem to be inappropriate. The view I will be taking of the moral agent is, in any case, one which is well enough entrenched in the literature to provide a plausible and interesting, even if contentious, point of departure.

In the remainder of this chapter I will deal with the various elements of the proffered sketch of the moral agent and I will argue that most of them can be generated, as it were, out of two fundamental features of the moral outlook: objectivism about questions of value and a concern for the interests of persons. I consider these features first. Finally, in the last section, I attempt a preliminary account of amoral agency in terms of these central features.

2 The objectivity of moral judgments

In this section I will argue that one central fact about moral language users is their supposition that moral judgments are true or false, or at least correct or incorrect, and that the standards for their assessment are objective and interpersonally valid. Later I will also argue that this objectivity explains to some extent other features of morality. The discussion here will

be brief since the question of objectivism is dealt with at some length in Chapter IV.

Grammatically speaking, moral language does not differ significantly from non-moral language. Indeed one usual way of demarcating moral discourse involves attending to certain words or phrases (e.g. "good" or "morally good," "ought" or "morally ought," "right" or "morally right" etc.). The presence of these terms in a sentence is then taken as a *prima facie* indication that the sentence belongs to moral discourse. Something like this seems necessary since there are no grammatical constructions peculiar to moral language on the basis of which to make the distinction. Moral judgments have the same form as ordinary statements of empirical fact. Even "ought" judgments whose function is quite distinctive can be transformed into ordinary empirical statements by substituting for "ought" a verb such as "forgot" or "decided" or "promised."

Not only do moral judgments have a grammatical form which makes the attribution of truth and falsity to them seem natural, they also function in ordinary language as if they actually do have a truth value. Unlike matters of taste, moral issues are typically ones over which the representatives of opposing views cannot content themselves with agreeing to differ, and this is not simply because moral issues have a practical urgency which matters of taste lack. Even when the parties to a moral debate have reached the point at which each has done all he can do to dissuade the others from their position, there is usually the residual belief among all that there is a "real" and correct moral solution to the problem.

Moral judgments, then, have a form and use which lend consid-

erable support to the idea that they, like other utterances of the assertive form, are either true or false. The moral agent accepts the idea which is implicit in his language that there are interpersonally valid standards for the assessment of actions, which do not refer to the wants, desires, preferences, etc. of the agent. But if moral judgments are true or false what makes them so, and how do we know which are true and which false? And if they are not statements of *some* kind of fact which can be true or false, what are they? The attempts to answer these questions form much of the history of moral philosophy in the Western tradition.

Ordinary language and ordinary moral thinking are, I think, saturated with the view that there is a "real" right and wrong, that there are objectively valid requirements on action which have a claim on persons regardless of their first-order interests, preferences and desires. This is not to suggest that there are no problems in this view; there are, and I will discuss them at length in Chapter IV. But the idea that morality, whether it is conceived of as a set of rules or principles or as a mode of practical deliberation, has an authority which resides in something objective and outside of the individual, seems to me an idea deeply engrained in the ordinary man's thinking. Moral judgments appear to most people to be true or false, in spite of epistemological problems of how we can know which are true and which false.

Obviously, not all practical judgments in which the terms "ought," "right," "wrong," "good" and so on appear are moral judgments and there is no need to suppose that the moral agent gives an objectivist interpretation to all of them. For example,

when we tell someone that he ought to see a certain play or that he is driving in the wrong gear, we are not typically invoking any standards of action which have an applicability beyond what they have as a result of a consideration of what we take that person's actual concerns to be. As a first approximation, let us take as paradigmatic of moral judgments, a judgment to the effect that some person ought to do something, where the speaker takes that judgment to be true or correct, and where the standard's applicability to the action in question is independent of the wants, desires, tastes and preferences of the persons involved. A person may, of course, want to do what he ought to do; indeed if he does the thing in question there must be *some* sense in which he wanted to do it but the force of the moral "ought" judgment is not dependent on his evident wants.

Closely connected with the idea that moral judgments can be true or false, or at least correct or incorrect according to interpersonally valid standards, is the view that moral judgments are overriding or supremely authoritative in practical matters. There is a real problem involved in understanding someone who sincerely believes that he ought morally to do something and yet fails to do it, since there seems to be nothing which could ultimately *justify* his action. In some sense, moral reasons are the very best.

There are at least two ways of accounting for the supremacy of the moral "ought," both of them consistent with the notion that "morally ought" entails "ought, all things considered." First, one might suppose that a moral judgment includes, by giving due and

proper weight to, all of the various sorts of reasons in a situation requiring a moral decision. This, what we might call the *summary view* of the moral "ought," is illustrated in the following passage by C.I. Lewis:

Other considerations notwithstanding, what is not *morally* right simply is not right to do We shall hardly conceive of the moral judgement as such a final arbitrament unless we also think of it as a judgement in which all valid claims upon the act in question are duly weighted and adjudicated. And the technical and prudential, as well as consideration of the consequences to others, are such valid claims.

(Lewis, 1969, *Values and Imperatives*, p. 7)

Secondly, it might be thought that moral requirements have a special character such that they are incommensurably weightier than claims of any other sort. If, for example, one thinks that moral questions arise only when the will of God is involved and that nothing could have any force in such situations except for consideration of what God wants, then the moral "ought" will be supreme because it has a special source and authority and not because it includes and balances a number of valid claims of various kinds.

I think that the summary view is the more plausible and it has the advantage of allowing for a rather broader range of moral problems. The summary view could be held in the extreme form in which any practical "ought" judgment counts as moral regardless of the nature of the situation in which it arose, but usually there are some constraints placed on the sort of situation in which a moral problem can arise. In the next section, I argue that the best candidate in this regard is the requirement that the interests or welfare of persons other than the agent be involved.

I said that the issue of the practical supremacy of moral judgments is connected with the thesis of the objectivity of moral judgments. Actually the connection is quite trivial since any summary "ought," being equivalent to "ought all things considered," will have the required authority. For example, in a situation in which the only valid reasons a person has are prudential ones, the judgment that he ought (all things considered) to do this or that will have a supremacy over any and all the *prima facie* "oughts" which go into the final judgment.

However, it is not very plausible to construe every "ought, all things considered," as a moral judgment. Even when the "ought" judgment involves accounting for the welfare of persons, it will seem like a moral judgment only if it can be supposed that there is a more or less determinate way of weighting the claims supposed to be *prima facie* valid which cannot legitimately vary depending on who makes the judgment or when it is made. In Chapter III I will examine some attempts to tackle the problem of showing that there are such valid modes of reasoning.

Most attempts at understanding morality have taken an objectivist line; however, there are metamoral theories which are non-objectivist and I will consider them later as regards their acceptability as analyses of moral phenomena. It will turn out that, to the extent that they are acceptable, they do not capture the practical thinking of an amoralist. For the present I must content myself with the unargued statement of the claim that metamoral theories which attempt to analyze moral judgments as expressions of emotion or attitude all fail to make sense of the life moral

concepts actually lead in ordinary thinking. In particular, they do not provide an account of the sense that most people have that morality involves real and objective claims, duties and obligations which are independent of their own personal emotions, valuations and attitudes.

Sidgwick, allowing for the emotional content of moral judgments says:

The peculiar emotion of moral approbation is, in my experience, inseparably bound up with the conviction, implicit or explicit, that the conduct approved is 'really' right - i.e. that it cannot, without error, be disapproved by any other mind (p. 27).

So far, then, from being prepared to admit that the proposition "X ought to be done" *merely* expresses the existence of a certain sentiment in myself or others, I find it strictly impossible so to regard my own moral judgments without eliminating from the concomitant sentiment the peculiar quality signified by the term 'moral.'

(Sidgwick, 1966, *The Methods of Ethics*, p. 28)

This seems to me to be a very good statement of the element of objectivism present in the ordinary moral consciousness.

A note about terminology: In the remainder of the thesis I will use the terms "objectivist" and "subjectivist" in two ways. Sometimes they will refer to a theoretical view about the nature of moral language, value judgments and so on. Sometimes, however, they will refer to the sort of person whose concepts and language are correctly described by these theories. If someone is highly reflective in his use of concepts he will hold a theory about the nature of those concepts as well as use them in the way described by the theory. A subjectivist, for example, may be someone who does not assign objective value to anything, does not recognize objectively valid

standards of action, etc. and who realizes that others do and consequently understands that his use of certain words may mislead others. I hope that which of these uses of "objectivist" and "subjectivist" is intended in various contexts will be clear, and that no serious confusions will result.

3 The welfare of others

The acceptance of the objectivity of judgments about what persons ought to do, or about what states of affairs have value and ought to be promoted, goes a long way, I think, toward constituting a person as a moral agent. But it certainly does not go all the way. Consider, for example, someone who believes that the only thing of real value is the existence of great works of art, or the existence of as much "unspoiled" nature as possible, or the supremacy of the Nazi state. Such a person can talk comfortably enough about what ought to be done, what obligations and duties people have, what actions are right and wrong, and so on; yet I think we would hesitate to call his position a moral one. For convenience, let us call a person who is prepared to make objective "ought" judgments of this sort, to advocate such values publicly and to defend them as valid and binding on all, in spite of the obvious conflict between these values and any true concern for the weal and woe of human beings generally, an ethical (but amoral) agent. An ethical amoral theory, then, will be a practical theory about how people ought to act, what obligations and duties they have and so on, which can be consistently maintained, acted upon and publicly advocated with some (perhaps only minimal) plausibility

and which claims to express the truth about what reasons people have for acting. To use Hare's useful terminology the ethical "ought," like the moral "ought," is universalizable and prescriptive (1963) and this commits those who use it to public discussion.

What I want to claim is that, as we normally understand the idea of morality, a moral agent must display some direct concern for the welfare of persons. Whether his concern rests on the belief that it is his duty to care for others or whether it rests on nothing beyond itself may not be crucial, but it cannot, for example, rest on a belief that his own happiness depends on making others happy. More about that later, but the basic idea is that there is a difference between taking pleasure in helping others and helping others in order to get a feeling of pleasure. Different moral theories will ground this concern for others differently but it makes little difference to my point here whether, for example, we are to respect others because they are endowed with the faculty of reason or because they are capable of feeling pain and pleasure.

W.D. Falk (1963, "Morality, Self, and Others") distinguishes what he calls the "formalist" and the "non-formalist" views of the nature of the moral "ought." According to the formalist, an "ought" judgment is normal if it expresses a person's overriding concerns or principles. On this view any resolution of a person's practical problems will count as moral provided that it embodies that person's highest (overriding) values. For the non-formalist, an "ought" judgment is moral only if it also involves a concern with the welfare of other persons. I have suggested that the latter approach is more in line with our ordinary understanding of moral

agency. Moral thinking, then, has a central formal feature - the objectivity of moral judgments and their attendant overridingness; and a central substantive feature - a manifest involvement with the welfare of persons.

Various philosophers have tried in various ways to account for the necessity of considering others in our practical deliberations. David Hume, for example, postulated a universal natural affection or sentiment directed toward the benefit of other persons and society in general. Only if men have such a sentiment, he thought, could they be moved to moral action. He says:

[T]hough this affection of humanity may not generally be esteemed so strong as vanity or ambition, yet, being common to all men, it can alone be the foundation of morals, or of any general system of blame or praise.

(Hume, 1966, *Enquiry Concerning the Principles of Morals*, IX, I, p. 111)

Hume's problem, however, was to show how this natural affect, supposing it exists, comes to appear in the guise of objective and external requirements on action - whence duties and obligations, etc.? Joseph Butler tried to solve this problem by postulating a faculty of conscience whose nature it was to issue authoritative directives on behalf of our various sentiments (notably benevolence) (Butler, 1967, esp. p. 53).

The central importance of the notion of the welfare of others to moral agency is further attested to by the fact that love of humanity, benevolence, kindness, brotherly love, charity and so on have always been primary moral virtues.

Not content with the postulation of a natural and universal sentiment as the basis for the requirement to display a regard for

others in moral thinking, some writers, R.M. Hare (1963) for example, have argued that the moral requirement to consider the interests of others is derivative from a formal or logical feature of moral discourse. Hare maintains that the distinctive features of moral language are its prescriptiveness and its universalizability, the latter being the source of the moral injunction to consider others. That is, any action which has consequences for persons other than the agent can be the subject of a moral judgment and since moral judgments are universalizable, the agent, if he concerns himself with the moral nature of his actions, must be prepared to accept (even prescribe) actions of the same sort by others in cases where the roles are reversed. This is because moral judgments imply the acceptance of a principle which applies to anyone. Thus, when considering what to do, the moral person considers the effects of his actions on others as though he were in their shoes, having their concerns. No general sentiment of benevolence is postulated as underlying this imagined reversal of roles; rather it is, according to Hare, a requirement of playing the morality game.

Other statements of such a Golden Rule approach abound. The idea is always basically the same; morality requires us to look to the effects of our actions on others as though we were in their position *because* of the logic of moral judgments, principles or rules. In spite of the wide acceptance of such arguments among philosophers, they fail to address the crucial questions, I think. It needs to be explained why anyone should be concerned to engage in moral deliberations and hence to subject themselves to the alleged logical implications of making moral judgments. Even if a deep

concern for others could explain participation in the moral realm, the use of moral language, etc., the very fact that moral requirements are so often seen as external impositions serving to modify the "natural" course of behaviour suggests this sentiment is not strong enough, in general and in itself, to account for the feeling that moral thinking is required. This may lead to the postulation of a distinctively moral sort of motivation such as Kantian "love of duty" or "reverence for the moral law" and to the idea that moral requirements are objectively valid and independent of even the limited concern which people have naturally for their fellows.

Setting aside for now the problems involved in attempting to show that a concern for others is somehow required, let me suggest what form this concern takes for the moral agent. The next section deals with the concept of the moral point of view, a point of view from which it is inevitable that one take full cognizance of other persons' interests. In Chapter III, I will take up the question of why it might be supposed that anyone is required to take up that special point of view in his practical deliberations.

4 The moral point of view

Moral thinking involves the adoption of a special point of view. Kurt Baier (1958) has termed this the "moral point of view" and the "God's eye point of view." The notion that there is some perspective above the concerns of the deliberating agent from which situations requiring a moral decision are to be viewed, however, goes back at least as far as Adam Smith (1976, *The Theory of Moral Senti-*

ments, esp. VII, ii, I.49). The basic idea is that to act as a moral agent one must, in deliberating on what to do, consider oneself as just one person among others and one must regard one's interests as just the interests of one person affected by the action of oneself considered as a "someone."

The concept of the moral point of view is connected with the idea that moral judgments are objective since deliberating from an "impersonal" point of view is helpful in rising above one's own particular inclinations to a view of one's situation in which one will not confuse one's own good (welfare) with what is truly (objectively, really, impersonally) good. I have discussed the objectivity of moral judgments briefly in a previous section and I will deal with it in more depth in a later chapter.

It has been argued that there is a point of view, that of an *ideal observer*, which is either involved in the very meaning of moral judgments or which helps to give the truth-conditions for moral judgments. Theories which hold that moral judgments can be analyzed as *meaning* that an ideal observer would do or think or feel in certain ways can be distinguished from theories which appeal to an ideal observer in order to clarify the nature of moral thinking. As an example of the latter, consider theories such as those advanced by moral sense theorists - the ideal observer here is usually some (non-actual) being endowed with an uncorrupted moral sense who is fully aware of the facts. Hume can be understood this way although he has also been taken to be an exponent of the former stronger thesis about the meaning of moral judgments (c.f. C.D. Broad, 1930, *Five Types of Ethical Theory*, pp. 84-93).

Roderick Firth (1951-2, "Ethical Absolutism and the Ideal Observer") more recently has attempted a revival of the ideal observer theory. His is a theory about the meaning of ethical judgments in which he proposes that we express the meaning of statements of the form "x is right" in terms of other statements which have the form, "Any ideal observer would react to x in such and such a way under such and such conditions" (p. 329). Firth's description of the conditions under which the ideal observer is supposed to operate are instructive, as is his discussion of the nature of the ideal observer's reaction. The ideal observer's reaction is to be one of the "specifically moral emotions" of moral approval or disapproval. Firth claims that there certainly seems to be such emotions which are experienced and therefore "unless apparent facts of this kind can be discounted by subtle phenomenological analysis, there is no epistemological objection to defining the ethically significant reactions of an ideal observer in terms of moral emotions" (p. 329).

The defining characteristics of the ideal observer are set out by Firth as including the conditions under which his reaction to situations are morally significant. He notes, incidentally, that:

[T]here is no good reason to believe that all human beings, no matter what the extent of their individual development, and no matter what their past social environment, could analyze their ethical statements correctly by reference to precisely the same kind of ideal observer and precisely the same psychological phenomena (p. 330).

With this caution in mind, Firth describes the ideal observer.

1. *"He is omniscient with respect to non-ethical facts."*

This reflects the fact that:

"We sometimes disqualify ourselves as judges of a particular ethical question on the ground

that we are not sufficiently familiar with the facts of the case, and we regard one person as a better moral judge than another if, other things being equal, the one has a larger amount of relevant factual knowledge than the other" (p. 333).

2. *"He is omniperceptive."*

This reflects the fact that:

"We sometimes disqualify ourselves as judges of certain ethical questions on the ground that we cannot satisfactorily imagine or visualize some of the relevant facts, and in general we regard one person as a better moral judge than another if, other things being equal, the one is better able to imagine or visualize the relevant facts" (p. 335).

3. *"He is disinterested."*

I.e., he is entirely lacking in particular interests where a particular interest is one which cannot be suitably described in terms of universals without altering the nature of the interest. This condition reflects the fact that:

"We sometimes disqualify ourselves as judges of certain ethical questions on the ground that we cannot make ourselves impartial ..." (p. 335).

4. *"He is dispassionate."*

This, along with the absence of particular interests, completes the analysis of impartiality mentioned in 3:

"... we can say that an ideal observer is dispassionate in the sense that he is incapable of experiencing emotions of this kind [particular emotions] - such as emotions of jealousy, self-love, personal hatred, and others which are directed towards particular individuals as such" (p. 340).

5. *"He is consistent."*

This characteristic is a consequence of all of the other characteristics and the presence or absence of consistency is a test of the adequacy of the characterization of the ideal observer. That is, the characterization of the ideal observer must guarantee that his reactions will be invariant among similar situations.

6. *"In other respects he is normal."*

The concept of normalcy is not easily analyzed but the idea is that the ideal observer is, after all, a person and his

personality should lie within the bounds of the normal (c.f. pp. 344-5).

I have quoted Firth at length because I think that he has isolated most of the factors which actually are used to discount certain persons as eligible for inclusion in serious moral debate and to discount particular moral judgments. Also his account gives a good indication of one way in which the interests and welfare of other persons may come to get a grip on the practical deliberations of the moral agent.

5 Some other features of morality and moral agency

So far I have argued that moral agency involves primarily the acceptance of the objectivity of moral judgments and a preparedness to take up a deliberative point of view from which the interests of other persons come to have considerable practical significance. In this section I wish to show that it is reasonable to take these two features to be central ones by considering their relationship to other aspects of morality and moral agency.

(i) Morality is practical - the problem of motivation

One thing which is utterly beyond dispute is the intimate connection between morality and action. As Hume put it:

If morality had no influence on human passions and actions, 'twere in vain to take such pains to inculcate it; and nothing wou'd be more fruitless than that multitude of rules and precepts, with which all moralists abound. Philosophy is commonly divided between *speculative* and *practical*; and as morality is always comprehended under the latter division, 'tis supposed to influence our passions and actions, and to go beyond the calm and indolent judgments of the understanding. And this is confirm'd

by common experience, which informs us,
that men are often govern'd by their duties,
and are deter'd from some actions by the
opinion of injustice, and impell'd to
others by that of obligation.

(Hume, 1888, *Treatise*, III, I, I,
p. 457)

While it is probably too much to say that every utterance properly classed as a moral judgment must be immediately relevant to action, there is an undeniable and essential connection between morality and questions about what to do. Clearly not all practical questions are candidates for moral solution and, perhaps less clearly, it is always possible to decide what to do without thinking morally.

At least this: anyone who sincerely makes a moral judgment thereby makes a commitment to act in accordance with that judgment insofar as he is capable of doing so, should the occasion arise. Moral judgments, in some sense, commit the will. R.M. Hare (1952, pp. 163-8) puts the matter in terms of moral judgments entailing imperatives. C.L. Stevenson (1944, pp. 13-19) speaks of moral judgments as expressing attitudes toward something, thereby ensuring the practical bearing of moral judgments and Kant made the good will the reference point of all moral evaluations.

W.D. Falk (1947-8, "'Ought' and Motivation ") introduced a distinction between two views on the relationship between moral "ought" judgments and motivation or between two different uses of morally "ought." On the *externalist* view, to say that someone ought to do something is to point to some objective or external requirement on him. The requirement may be external in the sense of being grounded in the commands of a deity, the dictates of

society or even in something as vague as the nature of his situation. The important thing is that the demand expressed in an external use of "ought" "has an objective existence of its own depending in no way on anything peculiar to the agent's psychological constitution" (pp. 125-126). The problem, Falk argues, with construing the moral "ought" as external is this:

If "I ought" means "I am from outside myself demanded to do some act," whether by the will of another, or more impersonally by the "situation," there will then be no necessary connection for anyone between having the duty and being under any manner of real compulsion to do the act. For no one really need do any act merely because it is demanded of him, whether by a deity or society or the "situation," but only if, in addition, he finds within himself a motive sufficient for satisfying the demand (p. 126).

But there is another use of "ought" besides the externalist one, according to Falk. It is what he calls the "internalist use" or "motivational use." When "ought" is used motivationally, to say that someone ought to do something is to:

express nothing other than a certain relation between a person's dispositional and occurrent motives: that though occurrently he had no impulse or desire to do an act or none sufficiently strong, dispositionally he was under an effective and over-riding compulsion to do it (p. 129).

In ordinary moral language, Falk maintains, these two uses are not adequately distinguished:

The external and internal uses of "ought" remain undifferentiated, and are imperceptibly juxtaposed and confused. There may be an unnoticed switch from the use of "ought" from one to the other, from a divine command, or requirement of others, to a dictate of conscience, or, when the more sophisticated speak the language of

objective "claims," an alternation between a nebulous externalist and an internalist interpretation of one and the same thing (p. 137).

There may well be cases in which "ought" is used externally and others in which it is used internally or motivationally. "You ought to drive on the right hand side," may be an example of the former and "You ought to watch what you eat," may be an example of the latter. However, it is not at all obvious that in moral discourse we need to make a choice between these two uses. An objectivist can, it seems, perfectly well argue that the moral "ought" is, strictly speaking, neither internal nor external, or that it is, in important ways, both. There is no obvious incoherence in arguing that moral "ought" judgments present demands which are independent of any particular person's peculiar psychological constitution and are to that extent external, but that they nevertheless recommend actions which anyone (or any fairly normal person about whom it is appropriate to make moral judgments) would be motivated to perform under certain conditions. Of course, those conditions would have to be grounded somehow in some very basic facts about human nature, rationality, the point of deliberation or something of that sort.

It should perhaps be noted that Falk is not very careful to distinguish "externalist" and "internalist" as applied to *uses* of "ought," to *senses* of "ought," and to *views* about the meaning of "ought." I think he supposes that there are at least the two uses of "ought" defined above and that there are at least two views about which use is involved in the use of "ought" in a moral context. According to Falk, we must decide which use or view best

reflects what we want to say using "ought" morally. I am suggesting that there is no need to choose between the internalist and externalist views about the nature of the moral "ought" since there may not be two distinguishable uses of "ought" in moral discourse.

I do not wish to suggest that there are no difficulties in the way of discovering an objectivist analysis of moral discourse which effects a synthesis between externalist and internalist views. I think there are very serious difficulties indeed as will become clear as we proceed; however, I disagree with Falk's claim that what one is dealing with in attempting to understand moral language is a confusion of two uses of "ought."

Certainly Falk is adequately aware of the appeal of objectivist analyses which attempt the synthesis. He notes:

People very commonly combine a view of "ought" as a requirement from outside, or an inner compulsion of a special quality, with adherence to a purist view of its connection with motivation, not so much as long as they view the moral law as the demand of a deity or of social convention, but once they think of it as somehow objectively grounded in the nature of things (p. 136).

I reject, then, Falk's claim that what is ultimately required is something in the nature of a decision between two uses of "ought." If the moral "ought" is to be understood at all, as it is used in ordinary moral discourse, it must be understood as an expression of an "external" demand or claim which does have, at least dispositionally, an intimate connection with the motivations of normal human agents. This must be so if, as I have argued, moral judgments are understood, by those who make them, as expressions of objectively valid requirements on action.

(ii) Morality as a social institution

Another feature of morality which virtually all metamoral theories recognize is that the use of moral concepts and moral thinking are essentially social. Richard Taylor (1970, pp. 125-128) makes the point that moral considerations are social in nature by asking us first to imagine a world devoid of human life and then one with a single inhabitant and finally one with one or more persons. Taylor argues, rightly I think, that no question of a moral kind can arise until the final stage at which we have persons interacting in a social manner. Recall that Kant formulates the Categorical Imperative in terms of rational agents living as a kingdom of ends.

Emphasizing the social character of morality, P.F. Strawson suggests a "minimal conception of morality" as a "public convenience." He says:

[I]t is a condition of the existence of any form of social organization, of any human community, that certain expectations of behaviour on the part of its members should be pretty regularly fulfilled: that some duties, one might say, should be performed, some obligations acknowledged, some rules observed.
(Strawson, 1961, "Social Morality and Individual Ideal," p. 5)

This conception is probably much more minimal than Strawson thinks unless some connection is made between "expectations," and "duties" and "obligations" since morality involves duties and obligations and not just expectations. It is a tautology that social organization involves persons fulfilling expectations "pretty regularly": still, every society does have a morality of some sort and that may at least be an interesting tautology.

The idea that morality has a *function* and that its function is the preservation of social life is a powerful one and linked to it are other features. For example, morality, as a system of rules and principles or as a mode of deliberation, must be teachable and must be such that human nature being what it is, people will for the most part adopt the rules or deliberative procedures and act in accordance with them. Morality cannot absolutely require what most people cannot be expected to do. The requirements of morality must not seem arbitrary or overly oppressive and they must carry authority of some sort in order that it be capable of curbing behaviour which would jeopardize the stability of the social unit. Most, if not all, of morality is thus concerned with modes of behaviour which potentially affect others. Even those areas of private morality (as opposed to social morality) which seem to concern no one but the individual concerned, can generally be understood in terms of the production of character types which are beneficial or at least benign from the point of view of the survival of society. If a boundary can be drawn between moral virtues and vices, and non-moral human excellences and deficiencies, it is likely to lie at the juncture of those things which are significant to the maintenance of societal stability and those which are not.

The social character of morality can be seen as a corollary of the fact that moral agents recognize the (objective) requirement to concern themselves with the welfare of others. Moral problems arise when the actions of one person or group of persons affect the welfare of others and this means that the appropriate

forms of social interaction (society's rules, conventions and laws) will be determined in large measure by the outcome of deliberations conducted in the moral mode. Where the actions of individuals do not affect others there is no need for anyone to consider anyone's welfare but his own, and there is thus no room for moral judgments. But in a social setting moral agents will set about determining (e.g. by pondering the situation or a range of possible situations from the moral point of view) how to achieve and enforce correct modes of conduct.

(iii) Morality adjudicates conflicts of interests

Connected with the idea that the moral agent considers the interests of others is the notion that it is of the essence of morality that it prevent as far as possible, and that it adjudicate where not possible, conflicts of interest among persons in a society. It is also often suggested that morality serves to remove conflicts of interests within a given individual (Plato and Butler come to mind). This view is put succinctly by G.W. Allport:

[A]ll theories of moral conduct have one primary purpose: they set before us some appropriate formula for handling conflicts - whether the conflicts be between warring interests in one individual or among individuals.

(Allport, 1959, "Normative Compatibility in the Light of Social Science," p. 139)

Many moral philosophers have held the view that morality can be considered as a device or method or set of principles whose primary function is the resolution of conflicts which arise from the fact that people often have desires or interests which cannot be all simultaneously satisfied. Kurt Baier goes so far as to say,

" ... by 'the moral point of view' we mean a point of view which is a court of appeal for conflicts of interest" (Baier, 1958, *The Moral Point of View*, p. 190). Often people who think of morality in these terms suppose that, morally speaking, all interests are equally legitimate taken by themselves; that is, if things were arranged in such a way that no one's interests conflicted with anyone else's interests there would be no need for - no reason for and no room for - morality. Many utilitarians, especially those influenced by modern welfare economists, tend to take this line. Other thinkers see the role of morality slightly differently, placing the emphasis on the prevention rather than on the resolution of conflicts. On this view morality concerns itself with the legitimacy of interests, evaluating them rather than taking them as givens to be dealt with according to some sort of calculus. I think the latter approach is closer to the way in which morality actually operates.

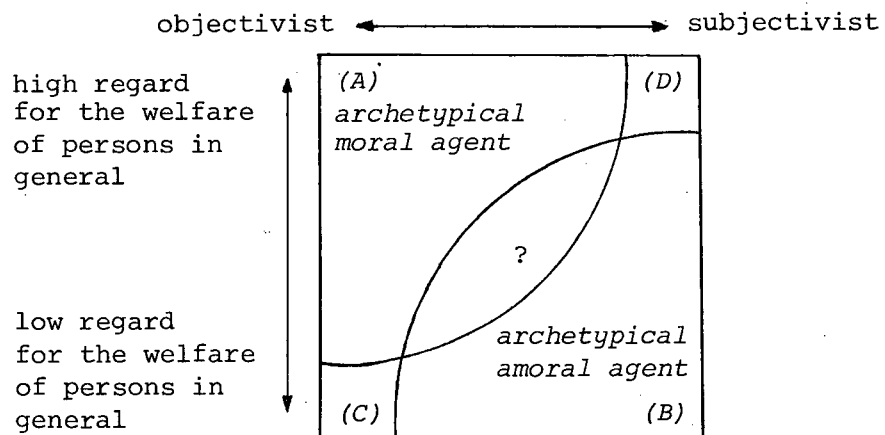
Given that moral agents concern themselves with the interests of others, it is easy to see how it comes about that morality should be seen as functioning to avoid or reconcile conflicts of interest. The idea that there is a correct resolution of situations in which individuals' interests conflict combined with a willingness to take everyone's welfare into account leads naturally to the view that not all interests are legitimate or that not all interests have a legitimate claim to be fully satisfied. The correct solution must be one in which everyone can do as he ought to do and a true concern for persons' welfare is consistent only with a solution which gives due consideration to the welfare (represented,

at least in part, by interests) of all. Since it is hard to see how everyone's welfare can be given due weight by a solution which recommends conflict and competition, we can expect the moral outcome to be one which specifies, in some way, which interests are to be satisfied to what degree. Thus morality's function of resolving conflicts is grounded in objectivism and concern for others which were earlier identified as the core of moral agency.

6 The amoral agent

If my characterization of the moral agent has been made sufficiently plausible, it should not be difficult to see that there are two main lines of enquiry open in our search for a way to understand amorality. First, if the acceptance of objectivism is a central characteristic of the moral agent, we can expect to find that a radical departure from this posture, toward subjectivism, will yield one form of amorality. Secondly, an absence of, or minimal, or highly selective, concern for the welfare of others in someone may also provide grounds for considering him to be amoral.

At the risk of being crude, the situation can be represented as follows:



Looking at things in this way we can isolate four extreme positions;

(A) The archetypical moral agent

This is the sort of person I have been describing in the present chapter and hopefully nothing further is required here by way of explanation.

(B) The archetypical amoral agent

To understand this sort of person properly would require that a great deal be said about subjectivism and since it will be more fully discussed later, I will simply sketch some main ideas here. The radical subjectivist denies that there is any very significant sense in the notion of things being "really" valuable or that anyone is subject to requirements or claims on his actions "from without." Things are valuable *only* in the sense that they are actually (or possibly potentially) valued by actual persons. The most extreme form of subjectivism which has received any very complete treatment is probably to be found in Sartre's *Being and Nothingness*. Sartre there advances the theory that in some very real and radical sense people do (or at least can) freely choose their values, their character and so on.

Because of his extreme subjectivism, the archetypical amoralist cannot make sincere moral judgments since to do so would involve him in the acceptance of interpersonally valid standards of action and would commit him to trying to justify his or others' actions on such grounds. Furthermore, the person at (B) takes little or no interest in the welfare of others and thus whatever explanations of his actions in "moral situations" he gave they could hardly be construed as contributions to a moral debate. He takes

no pains to adopt the moral point of view since to do so would be to remove himself from any concern with the world he might have had. These claims will, I hope, become clearer through the discussions which appear in Chapter III section 6 and following.

(C) The objectivist amoralist

Consider someone who is fanatically committed to the preservation of wildlife to the total neglect of the welfare of anyone. He may even be prepared to sacrifice (or merely give) his life for the cause. Not only is he committed to this goal, he regards it as one which everyone ought to embrace. It just is very important (objectively speaking), according to him, that the wildlife of the planet be preserved. It is not (anyway primarily) that he wants to have the animals preserved for future generations, nor that he believes that human souls inhabit the bodies of the birds and beasts, let us say. Rather, for whatever reason, (divine revelation?) he thinks it important to follow the course he recommends. Other cases can be manufactured as well, perhaps less fantastic ones.

Actually there is one sort of objective amoralist which has received considerable attention - the egoist. The next chapter is devoted to a discussion of egoism.

In my sketch I have indicated that it might be possible for someone at (C) to be considered a moral agent. I am not very happy with this idea for reasons I have already given, but some philosophers seem willing to embrace the possibility less reluctantly. Hare, for example, leaning heavily on the fact that the person at (C) can use a practical vocabulary in which all of the logical structure of moral discourse as he defines it is present, seems

tempted to grant the fanatic (his term as well) the status of moral agent (my term). Ultimately, however, he says that "it is characteristic of moral thought in general to accord equal weight to the interest of all persons" (Hare, 1963, *Freedom and Reason*, p. 177).

(D) The other-regarding subjectivist

In many ways the person at (D) is the most interesting and I deal with the complexities of this case in the last chapter. It seems to me that there is here some possibility of moral agency of a special sort but further discussion must wait. Whether an other-regarding subjectivist is to be classed as an amoral or as a moral agent may depend on such things as how stable his concern for others is and what form his concern takes conceptually for him.

II

EGOISM

1 Egoism as a form of amoralism

The moral agent recognizes, in the interests of other persons, a source of valid claims on his actions. These claims may arise directly from a belief that others' welfare has an objective value or indirectly from an acceptance of a requirement that he deliberate from the moral point of view, taking no special cognizance of his own particular valuations and interests. If a recognition of claims grounded in the welfare of others is, as I have argued, a necessary component of moral thinking, someone who does not recognize such claims will be an amoral agent. There may be people who take minimal interest in anyone's interests, including their own; e.g. nihilists with little interest in anything, idealists, fanatics, religious zealots, etc. who care only for the greater glory of the type man, the supremacy of the Nazi state, the devotion of the human race to the Almighty, and so on. Such people are rare, but there is a type of person, the egoist, who strikes one as being rather less bizarre. The egoist, at least, cares about his own welfare, as we all do, and his peculiarity lies only in his disregarding the welfare of others, except when their welfare is of more or less direct consequence to his own.

Egoism, like the forms of "fanaticism" mentioned above may

be held objectively or subjectively. The objectivist egoist thinks that the only real reasons anyone has or could have for doing some action is that doing so conduces to his own welfare. This sort of view is usually called ethical egoism and is expressed in the claim that everyone ought to pursue his own interests. Egoism may also be held subjectively, however. The subjectivist egoist need not make any general claims about who has what sorts of reasons - he simply does care only about his own weal. If he holds his position reflectively we can expect that he will have some views about the nature of practical reasoning. But if he is truly a subjectivist he will not suppose that any error is necessarily being made by non-egoists, since there are no objective values (whether relativized to individuals or not). In the next section I consider some views which have been advanced on the relationship between morality and self-interest. Later I will consider in more detail some various forms of egoism.

2 Morality and self-interest

Because all of us are considerably concerned about ourselves and are sometimes tempted to ignore the interests of others, egoism has a certain interest and appeal. Yet most of us suspect that there is something fundamentally wrong (and not just from a moral perspective) with egoism. The problem is expressed by Sidgwick in his *Methods of Ethics* in this way:

I find that I undoubtedly seem to perceive, as clearly and certainly as I see any axiom in Arithmetic or Geometry, that it is 'right' and 'reasonable' for me to treat others as I should think that I myself ought to be treated under similar conditions,

and to do what I believe to be ultimately conducive to universal Good or Happiness (p. 507).

[E]ven if a man admits the self-evidence of the principle of Rational Benevolence, he may still hold that his own happiness is an end which it is irrational for him to sacrifice to any other; and that therefore a harmony between the maxim of Prudence and the maxim of Rational Benevolence must be somehow demonstrated, if morality is to be made completely rational. This latter view, indeed ... appears to me, on the whole, the view of Common Sense: and it is that which I myself hold (p. 498).

(Sidgwick, 1966)

Many efforts have been made to reconcile these conflicting intuitions by attempting to show that the conscientious moral agent will actually best serve his own interests. If it could be shown that morality and self-interest coincide or actually amount to the same thing, in some sense, then Sidgwick's problem would evaporate.

Plato sought to show that if we only understand happiness properly, we will be able to see that only the just (the moral) man can attain that true happiness which depends on a well-ordered soul. Even if one could accept Plato's psychological views, it is difficult to see why precisely those actions produced by moral deliberation should everywhere and for everyone be just the ones which conduce to the production of, or flow naturally from, the presence of an inner harmony, and it is even more problematic how a notion of moral obligation (which we suppose can override prudence) is possible, let alone necessary, on such a view.

Joseph Butler thought that the crucial distinction is not the one between self-interest and morality, but rather the one between those passions directed at one's own welfare and those

directed at the welfare of others, severally or as mankind. More importantly, Butler thought that there is in each of us a faculty which arbitrates among these passions, bestowing on some on various occasions a certain authority. This faculty, which he calls "Conscience," seems to be a cognitive faculty as well as being the distinctively moral one. Passions have many sorts of objects, the two most important of which are one's own welfare and the welfare of others. But both of these sorts of passions *in fact* can be satisfied by the same line of conduct. Conscience, as a cognitive guide, reveals the course of action which will lead to the satisfaction of both kinds of passion, whether we realize it or not. In Butler's words:

I must however remind you that though benevolence and self-love are different; though the former tends most directly to public good, and the latter to private: yet they are so perfectly coincident, that the greatest satisfactions to ourselves depend upon our having benevolence in a due degree; and that self-love is one chief security of our right behaviour towards society.

(Butler, 1967, "Upon Human Nature,"
I, *Fifteen Sermons*, pp. 35-36)

Butler's theory partakes of all the difficulties of intuitionist theories and one suspects that Butler's faith in the ultimate congruity of benevolence and self-interest should have given him more trouble than it apparently did.

Butler and Plato argue for a virtual identity of self-interest and morality, but they both make an appeal to some notion of a person's "real" self-interest in order to explain the apparently obvious fact that self-interest and morality *do*, at least sometimes, conflict. Plato's idea of self-interest rests on his concept of true

happiness which involves the notion of the natural supremacy of the faculty of reason; reason serves to keep the potentially unruly and disruptive passions in check, uncontrolled passion being the source of both evil and happiness. Butler sees our true happiness as consisting in the fulfillment of our social nature as well as our private nature and it is the function of conscience to show the way to this dual satisfaction.

It is probably fair to say that insofar as attempts to show the coincidence of morality and self-interest do justice to our ordinary conception of morality, they distort our ordinary notion of self-interest, "moralizing" the latter in a way which provides the desired conclusion (that it is in our "true self-interest" to be moral while leaving us with the original problem in a new form. Why follow our "true" instead of our "apparent" self-interest? On the other hand, insofar as such efforts do justice to our ordinary concept of self-interest, they typically fail to show that there are no real and troublesome conflicts between a person's self-interest and the requirements of morality.

The point of view of self-interest and the moral point of view are distinct. This would be true even if it could be shown (as I think it plainly cannot) that it makes no difference *practically* which point of view one adopts. Even if it were the case, that deliberation from either point of view would always, if done correctly, recommend the same actions, there would be significant differences between the moral agent and the egoist. The attitudes, values, and general temperament of these two sorts of persons, as well as their susceptibility to various sorts of considerations

and argument forms, would differ radically.

Kai Nielsen (1963) adduces "weighty" considerations of a mundane sort in favour of the individual's taking the moral point of view." The considerations he offers are designed primarily to show that acting in certain "immoral" ways is likely to produce psychological strains and imbalances which tend to produce personal unhappiness. The moral teachings of society are lodged deep in our subconscious and are not easily ignored or overridden. Acting "immorally" will also often make interpersonal dealings very difficult. But, however successful these arguments are in getting the egoist to act morally, it is doubtful that they can go very far toward erasing the difference between the egoist and the moral agent. For one thing, the egoist's conformity to the requirements of morality and even his adoption of the moral point of view in deliberation are mere expedients. They represent no appropriate recognition of the validity of the claims of the interests of others on him. Of course, it might be held that the "mundane considerations" Nielsen offers will suffice to get the egoist involved in (insincere) moral thinking and that from there something will happen to effect a genuine conversion to moral agency, but that is another story.

3 The concept of self-interest

I have argued that many attempts to minimize the conflict between morality and self-interest run afoul of our ordinary understanding of what the latter involves and something further should be said in this connection.

To begin with, it would not do to analyze the concept of self-interest in a way which destroys the distinction between the idea of someone's acting from self-interest, and the more general idea of someone's acting to produce some state of affairs which he is interested in bringing about. This would not yield a view properly called "egoism" at all since it would make the most dedicated moral agent out to be concerned with his own self-interest. That is, while there is some sense in which the dedicated moral agent is interested in doing the morally right thing, it would hardly accord with common usage to say that he is acting from self-interest. "Self-interest" is a concept related to interest *in* oneself and not to interest *of* oneself. Our normal concept of self-interest has *content* insofar as there are certain sorts of things which are taken to be in a person's interest regardless (in large measure, at least) of what he takes an interest in, although there is undoubtedly a component of self-interest which depends on what concerns a person actually has.

It is not difficult to begin a list of things which have a place here: health, wealth, happiness, respect, intelligence, opportunity, security and so on are prime candidates. I am not interested in completing or refining such a list - the concept is not likely determinate enough for this anyway - but it is clear that it could only be termed miraculous if it should turn out that the maximal satisfaction of all items on the list for everyone should turn out to require just those actions which deliberation from the moral point of view would require, and that is enough to show that morality is not straightforwardly a matter of self-

interest. Insofar as we are considering egoism as a form of amorality in virtue of the fact that the egoist has little or no concern for the welfare of others, it is not really very important exactly how we analyze self-interest. Any very plausible account of our ordinary use of the concept will sustain the necessary distinction between the egoist and the moral agent.

4 Ethical egoism - an objectivist amorality

The most commonly discussed formulation of egoism is the position which has come to be known as ethical egoism and which is usually stated in the form, "Everyone ought to pursue his own self-interest." The view certainly has the appearance of an ethical theory, but since anyone who holds it will not concern himself to any extent with the welfare of others it cannot be considered a moral theory.

Recall that in the last chapter I defined an ethical theory as a general objectivist practical theory about what sorts of reasons people have, what sorts of things are valuable, what sorts of states of affairs ought to obtain, and so on. The theory, to count as ethical, must be one which can be publicly defended. Again, to use R.M. Hare's terms, an ethical "ought" is prescriptive as well as universalizable. It is a theory about how people ought to act and the "ought" in its formulation must "carry the will" so that someone who holds the theory wants, in some sense, people to do as they ought to do. For most "fanatics" this is not a serious problem - the Nazi, for example, can wholeheartedly hope that everyone will contribute in whatever way they can to establishing the supremacy

of the Third Reich and he can publicly advocate his views. The ethical egoist, however, has problems here. Indeed there is some question whether ethical egoism is a *tenable* theory. A good deal of attention has been paid this question in recent years and it is worth pausing to consider the issue. Many arguments have been advanced to the effect that ethical egoism is either not a theory of *egoism* at all or else it is not an *ethical* theory or else that it is not a *consistent* practical theory. The main arguments can be summarized as follows:

1. If the ethical egoist's doctrine is an ethical one (a quasi-moral one) he must be able to promulgate it publicly; that is, he must be able to advocate that others accept it, pursuing their own interests even when their interests conflict with his own. But, since it is almost certain that conflict will occur, this would be to act against his interests. Therefore, the ethical egoist cannot both accept the doctrine himself and advocate the universal acceptance of it. Thus, anyone who adopts the doctrine as an ethical theory is involved in a practical inconsistency.
2. If ethical egoism is an ethical theory then the "ought" in its formulation must "carry the will" in the sense that the ethical egoist must will the state of affairs in which everyone pursues his own interests. But then his own interests are not his main concern and he is not an egoist at all. (See Campbell, 1972, "A Short Refutation of Ethical Egoism.")

3. It would be irrational or odd in some serious way to will that someone pursue some end and at the same time remain indifferent to, or even opposed to, his achieving that end. But the ethical egoist is in just this situation when his interests clash with others'.
(See Silverstein, 1968, "Universalizability and Egoism"; Medlin, 1957, "Ultimate Principles and Ethical Egoism"; Narveson, 1967, *Morality and Utility*, pp. 268-271.)
4. If someone ought to do something, then no one ought to attempt to prevent him from doing it. The ethical egoist acts and counsels others to act in ways which are inconsistent with this analytic truth to which he is committed by his use of "ought" in its ethical sense in his doctrine. Thus the ethical egoist is not able to accept the principle he thinks he accepts.
(See Baumer, 1967, "Indefensible Impersonal Egoism"; Baier, 1973, "Ethical Egoism and Interpersonal Compatibility.")
5. The ethical "ought" commits one to the view that what satisfies the "ought" is objectively good and ought to be promoted. The ethical egoist is committed to the view that there are many incompatible states of affairs which ought to be promoted, one for each person roughly. (See Moore, 1903, *Principia Ethica*, pp. 96-105; Quinn, 1974, "Egoism as an Ethical System.")

Other arguments can be found, for example, Glasgow (1968, "The

Contradiction in Ethical Egoism"); Singer (1959, "On Duties to Oneself").

Some attempts have been made to save ethical egoism from these problems. Ayn Rand (1964) and members of her Objectivist following (eg. Branden, 1970) defend a form of ethical egoism which avoids the arguments. But they succeed only to the extent that they are prepared to defend a very special and not at all ordinary conception of a person's self-interest, according to which it is in the person's interests to be a certain sort of ideal rational agent which among other things involves a commitment to honesty, integrity, and so on.

John Hospers has suggested that the objections above can be avoided by supposing the egoist to be saying something like this:

"I hope each of you [everyone] *tries* to come out on top," or "Each of you should *try* to come out the victor." There is surely no inconsistency here. The hope he is expressing here is the kind of hope that the interested but impartial spectator expresses at a game. Perhaps the egoist likes to live life in a dangerous cutthroat manner, unwilling to help others in need but not desiring others to help him either. He wants life to be spicy and dangerous; to him the whole world is one vast egoistic game, and living life accordingly is the way to make it interesting and exciting.

(Hospers, 1961, "Baier and Medlin on Ethical Egoism," p. 16)

Under such a formulation, the doctrine retains some of its egoistic flavour, although it is clear that the thing of greatest interest to the egoist is no longer his own welfare *simpliciter* but his attempting to achieve his self-interest in a world in which others are attempting to do the same, or perhaps it is simply

to live in a world in which everyone attempts to do as well for himself as possible and where the actual achievement of his goal in this game has no more intrinsic interest for him than the protection of his king in a game of chess. This is a possible theory, I suppose, but it is a peculiar one since, if the "ought" is the normal one of ethical theories, the egoist must be taking an odd disinterested view of his interest or else the satisfaction of his interests has value for him only under a condition of competition and not in itself. This makes the theory look less egoistic than it might at first, but it must be admitted that an ethical theory of this sort could be held. It is not a moral theory because there is no provision for a settlement of conflicts of interests and that is one mark of a moral theory. The theory does not hold that the strongest should win; the actual outcome of any struggle is ethically irrelevant although each egoistic participant can be expected to prefer winning to losing, as in any game.

Attempts to preserve the cogency of ethical egoism as a truly *ethical* theory, i.e. as a theory which can be articulated and advocated as an alternative to morality (in the sense of serving as a universal theory about how people ought to act without necessarily satisfying all of the constraints on a moral theory), seem ultimately very desperate. They also, through their efforts to preserve the "ethicalness" of the theory, inevitably fail to capture the egoism of the original position. Of course it can hardly be considered a fatal objection to a theory that it would be misleading to call it a kind of egoism. The only reason I

have drawn attention to the fact that some attempts to save ethical egoism turn out to involve the defense of some non-egoistic theory is that such alternative theories usually must forgo most of the intuitive plausibility that egoism gains by its claiming the support of the generally recognized rationality of acting in one's own interest.

5 Personal egoism - another objectivist amoralism

I have been discussing egoism as an ethical theory, that is, as a theory someone might hold about how people ought to act or about what sorts of reasons people have, where the force of "ethical" has been to retain the requirements that the theory be one which is objectivist and can be advertised and promulgated as a universally adoptable practical theory. This much an ethical theory shares with a moral theory and thus I have used the term "quasi-moral."

There is a tendency to analyze a statement to the effect that someone ought to do something as roughly equivalent to a statement to the effect that the person in question has (conclusive) reasons to do that thing. This tendency is harmless in some contexts but it is to be resisted in others, I think. I have used the formulation of ethical egoism in terms of what people ought to do and I have associated the term "ethical" with a specific interpretation of the force of the "ought" in that judgment. In particular, I have supposed that to make such an "ought" judgment is to commit oneself, in some sense, to wanting or willing that people act in accordance with it. There is something peculiar, if

not actually self-contradictory, in saying, for example, "Jones ought to do that, but I hope he doesn't." Expressed ought judgments imply that the speaker is prepared to support, encourage, and promote action in accordance with those judgments. Statements about what reasons someone has, on the other hand, do not carry the same sort of implications. "Jones has every reason to do it but I hope he doesn't," is not an odd thing to say.

Thus, egoism stated in terms of reasons, (still, remember, as an objectivist position) does not have quite the same flavour. The view that the only or the very best reasons anyone can have are reasons of self-interest, does not seem to require any commitment of the will where the actions of others are concerned. This view I will call personal egoism because someone can hold it and yet be under no requirement to advocate it or even to admit he holds it. The ethical egoist and the personal egoist agree about the nature of practical reasons but the ethical egoist supposes (mistakenly I have argued) that his position can be consistently promulgated. The ethical egoist's conception of, and attitude toward, his practical theory is one which cannot be maintained on reflection and so he must alter his conception of self-interest (in effect exchanging it for something else) or fall into personal egoism. It would be misleading to say that personal egoists hold the same practical theory as ethical egoists and differ only in their strategy. The ethical egoist thinks people *ought* to act in a certain way and that commits him to more than the personal egoist is committed to.

Henry Jack (1969, p. 479) suggests that egoism might be

understood as an "attempt to state the truth about ethics." Egoism, on his interpretation, is a philosophical theory about the reasons people really have and he admits that it could be promulgated consistently only by someone whose greatest happiness lies in the promulgation of philosophical truth. If egoism does state the truth about ethics there is no evident reason why anyone must feel bound to advertise that truth or defend it publicly.

The personal egoist need not recognize any obligation to tell the truth and, in the pursuit of his selfish ends, he may well find it expedient to withhold what he takes to be the truth about what reasons people have and even to lie about this. He acts so as to serve, as best he can, his own interests and when he believes that others have reasons of a similar kind to act contrary to his interests, he generally makes no effort to tell them so. Because he recognizes no requirement to tell the truth he sees no reason in principle, to tell the truth about ethics. It is interesting to note that personal egoism is not the only practical theory which opens the possibility that concealing the truth about ethics may be required. Plato can be read as endorsing this sort of deception, albeit toward a different end, and any teleological theory which makes the achievement of some end state the appropriate goal of action must reckon with the possibility that achievement of that end may require that people generally not be told that they ought to aim at it.

The personal egoist, of course, is not open to the sorts of argument we encountered above against ethical egoism since they rest on the fact that the ethical egoist is required, by his use

of the ethical "ought," to promulgate his view. But does this mean that no attack is possible against the position of the personal egoist? I think there are at least two ways in which he can be confronted.

First, since the personal egoist thinks that his view of ethics is the true one, we could, if we knew more about his reasons for thinking that people only have reasons of self-interest, open debate on that front. It is not inconceivable that he could be brought to admit that he was mistaken in this regard. Secondly (although if the project is to get the personal egoist to change his behaviour, it is not clear how this would help), one could argue that he is mistaken in thinking that there is *any* truth about ethics in the way he supposes. That is, there may be a way of showing the personal egoist that he is mistaken in thinking that there is any objective validity to the reasons he (or someone else) has for acting. This line of argument seems to me to undermine morality as well as personal egoism at least if my characterization of morality as it is ordinarily conceived was correct, but it does open a line of investigation. Once the egoist gives up his objectivism it may be possible to get him to change his practical basis so as to include a concern for others. As long as someone thinks that his *real* reasons are all ones of self-interest he may not be very susceptible to certain modes of thinking. Whether we count it as a form of argument or not, rubbing someone's nose in the facts of human existence can be effective.

Since I am here interested in the personal egoist who accepts the notion of ethical truth, I shall shelve the second of the above

approaches temporarily and concentrate on the first. What sorts of arguments then are available against objective personal egoism? The most promising and most popular approach is one according to which there are actually good reasons of self-interest for being moral; that is, it locates the personal egoist's mistake in his failing to fully appreciate the nature of the reasons he already recognizes. (In the next two sections I will be considering some arguments which have been recently advanced against egoism. While most of these discussions could be applied to ethical egoism, I take it that enough has been said in the previous section to discredit that view and consequently I shall assume that the personal egoist is the target of what follows.)

6 Kurt Baier

Since Socrates tried to satisfy Thrasymachus that justice profits the just man, philosophers have tried time and again to show that somehow morality can be grounded in self-interest. One of the best known recent attempts is Kurt Baier's *Moral Point of View* (1958) in which it is argued that there are "the very best reasons" for taking up the moral point of view.

Baier states his case in terms of types of reasons and tries to show that moral reasons are better than, or superior to, reasons of every other type. He begins with the claim: "The very *raison d'etre* of a morality is to yield reasons which overrule the reasons of self-interest in those cases when everyone's following self-interest would be harmful to everyone" (p. 309). Since everyone would be better off in a world where everyone follows moral rules

than in a world where everyone follows self-interest, Baier argues, moral rules are *superior* to rules of self-interest. But once we see that moral rules are superior to rules of self-interest, we have been given all the reason we need or could want for following moral rules. (Actually Baier talks as though he has established the superiority of moral rules to *all* other sorts of rules.)

The argument rests crucially on the notion of one kind of reason's being superior to another. Aside from the question of why we should suppose that all reasons of a given kind must be superior to all reasons of another kind or vice versa, and aside from the question of why the reasons which are superior for one person should be thought to be superior also for everyone else, there is a problem of why we should feel inclined to accept Baier's criterion for superiority as the relevant one. It would seem that if it makes sense to speak of choosing what sort of reasons are to be treated as superior, then the relevant criteria should come from a consideration of the point of deliberation (i.e. of using reasons at all). Baier himself, tries to do just this:

Our very purpose in "playing the reasoning game" is to maximize satisfactions and minimize frustrations. Deliberately to frustrate ourselves and to minimize satisfaction would certainly be to go counter to the very purpose for which we deliberate and weigh the pros and cons. These criteria are, therefore, necessarily linked with the very purpose of the activity of reasoning (pp. 301-2).

If we grant Baier's claim that we "play the reasoning game" in order to maximize our satisfactions and minimize our frustrations, what we need is some link between this and the concept of the superiority of one sort of reasons over another. His attempt to

provide this link has already been noted: reasons of type A are superior to reasons of type B just in case everyone would be better off (experience more satisfactions and fewer frustrations) in a world wherein everyone gave A-reasons priority over B-reasons. The trouble here is that while it may be up to an individual to treat moral reasons as superior (to deliberate from the moral point of view), it is not up to an individual to have everyone do so. If in choosing what sort of reason to count as superior, one were faced with the choice between a world in which everyone follows moral rules and a world in which everyone follows some other sort of rules, then one *would* have a reason (related to one's concern to maximize one's satisfactions) to opt for morality. However, this is clearly *not* a choice which does confront anyone.

Baier's argument could be made good if he could establish the (obviously false) premise that it is not possible for anyone (the egoist in particular) to maximize his own satisfactions without following reasons of the "superior" kind. But the egoist may well be in a position such that he is best off following reasons of self-interest while others also pursue his interests, second best off if he follows reasons of self-interest while others follow moral rules, third best off if he follows self-interest and others do whatever they will in fact do (assuming, as it is reasonable to assume, that his choice will not greatly affect others' choices), fourth best off if he and everyone else follow self-interest, and worst off if he follows morality and everyone else follows self-interest. If the point of deliberation is to maximize satisfactions then the egoist is best off following reasons of self-interest

it seems.

The only way I can see of making Baier's argument cogent is to suppose that it is really premised on the idea that the point of deliberation is to generate rational action and that to act rationally is to follow the best reasons and that the best reasons are the ones identified by his criterion of superiority. I think this probably is the way Baier intends the argument to run, given his talk of superior sorts of reasons. Still it is paradoxical to attempt to ground the superiority of moral reasons in an appeal to the self-interested point of deliberation in the way he does.

Indeed, Baier notes the circularity of trying to give moral reasons for adopting the moral point of view and the impossibility of giving ordinary reasons of self-interest for taking up a point of view designed to override reasons of self-interest. But his attempt to give the sense in which moral reasons are supposed to be superior to, yet somehow based on, self-interest is not very satisfactory. If the reasons for taking up the moral point of view are neither moral nor self-interested ones, one wants to ask, with Baier, "And what other reasons are there?" He answers:

The answer is that we are now looking at the world from the point of view of *anyone*. We are not examining particular courses of action before this or that person; we are examining two alternative worlds, one in which moral reasons are always treated by everyone as superior to reasons of self-interest and one in which the reverse is the practice. And we can see that the first world is the better world, because we can see that the second world would be the sort which Hobbes describes as the state of nature (p. 310).

The problem with this sort of appeal is that it is ineffective when addressed to the personal egoist, the sort of person it is designed to convince. The point of view of "anyone" is, essentially, the moral point of view and it does no good to try to convince someone to take up the moral point of view by pointing out that from the moral point of view this or that will appear to be the case. The personal egoist sees things from a particular point of view - his own - and unless there is something demonstrably wrong with his having that point of view, Baier's approach must appear quite irrelevant.

David Gauthier (1967), in his excellent article "Morality and Advantage," argues that there is an important distinction between the question "Why should I be moral?" and the question "Why should we be moral?" The latter question is one which can be addressed by the sort of considerations advanced by Baier, at least in part. If it were a complete characterization of morality to note that it is a system of principles or rules such that everyone is or can expect to be better off if that system rather than any other system were followed universally, then it is certainly plausible to argue that we, as some sort of deliberative community, have reason to adopt morality. Gauthier suggests that such a characterization is not a complete account of what we normally understand morality to involve and I think he is right in this. But beyond that the problem remains, as he points out, that the egoist does not deliberate as though he were choosing for such a community. What is needed it seems, is an argument to show that there is something non-morally amiss with the egoist's position. The argu-

ments in the next section could be directed at both the personal egoist and the subjectivist egoist considered in section 8, but, since I think they do no significant damage to either position, I will deal with them as if they were directed just against the personal egoist.

7 David Gauthier

It is important to note at the outset that David Gauthier, in the articles considered in this section, directs his attention to a class of agents which is considerably broader than the one I have been dealing with. In particular he is concerned with persons whose values are not confined to those "selfish" ends I allow the egoist. In "The Impossibility of Rational Egoism," (1974) he says: "An egoist is a person who on every occasion and in every respect acts to bring about as much as possible of what he values" (p. 442). Similarly, in "Reason and Maximization" (1975), he considers a person who always acts so that "the expected outcome of his action affords him a utility at least as great as that of the expected outcome of any action possible for him in the situation" (p. 418). The values or utilities of the agent Gauthier considers may concern virtually anything, including the welfare of others. I think it is inappropriate to call such a person an egoist, but that is not important for the assessment of Gauthier's arguments since my egoist is a special case of his egoist. If his arguments succeed in the more general case they succeed in the more particular case I have concentrated on. I propose to show that the arguments fail in the general case and there is

no reason to suppose similar arguments would succeed against the narrower version of egoism described above.

(i) Incompleteness

First, consider Gauthier's attempt in "The Impossibility of Rational Egoism" to show that there is a serious internal problem in the theory of egoism. Gauthier stresses the difference between the policy of trying to do as well for oneself as possible understood as an "over-all concern to maximize one's utilities" and that policy understood as the attempt to act *on each and every occasion* in a way which will maximize one's utilities. The former policy is consistent, for example, with a person's making promises and keeping them without raising the question of whether or not the promised action is to his benefit when the occasion for performing it arises. It is, then, only to the policy of always looking to utilities on all occasions that Gauthier addresses himself. The argument goes as follows. There are situations in which egoism does not uniquely identify a single action or strategy as the one which will maximize the egoist's expected utilities. Therefore, egoism is not a complete practical theory. Therefore, "practical rationality cannot be identified with unrestricted maximization" (pp. 455-6).

I am not as interested in the details of the situation Gauthier constructs to show his incompleteness result as in the underlying idea that a practical theory which is incomplete is unacceptable. Put another way, my concern is with the claim that there must be a unique rational action in any situation. I think it is clear enough that Gauthier is committed to this claim. He

says:

Before one can measure egoism against an external standard, one must measure it against an internal standard. One must decide what conditions a principle of action must meet, to be an *egoistic* principle, and then ask whether any principle of action can actually meet them. If no principle can meet these conditions, then egoism lacks internal rationality.... If egoism collapses into internal incoherence, rational egoism is impossible.

I shall demonstrate that no complete principle of action meets the conditions of egoism. By a *complete principle of action* I mean a function whose domain includes every possible situation in which a person might find himself and whose values include every possible action he might perform. A complete principle of action determines an action for every possible situation. I shall show then that the conditions of egoism cannot be satisfied by any function defined over all possible situations in which one might act and specifying an action for each of those situations (pp. 440-1).

If I understand this correctly, it means that any principle of action which fails to specify an action for every possible situation in which an agent finds himself, is incomplete and therefore "lacks internal rationality" and "collapses into internal incoherence." Further evidence that this is Gauthier's position is provided in his concluding section where he says:

Furthermore, if it is not possible always to act egoistically, then practical rationality cannot be identified with unrestricted maximization. And one may suppose that, if egoistic behaviour is ever rational, it must be shown to be rational by derivation from some more general characterization of practical rationality, which can be applied to all situations (pp. 455-6).

Now, how plausible is this completeness requirement? I should like to argue that it is not really very plausible at all.

In the first place, it seems that there are a number of kinds of situation in which no principle of action can get a grip. Consider, for example, a situation in which nothing of consequence is at stake, or in which there is nothing to choose among the alternatives - all of them being equally desirable or undesirable. To require a principle of action to pick out one of the possible actions in these circumstances is very peculiar. Or suppose someone finds himself in a situation in which he has virtually no reliable information on what the effects of his actions will be. Here again the failure of an otherwise plausible practical theory to recommend any particular action seems more a strength of the theory than a weakness. In short, it appears that any acceptable principle of action *ought* to remain silent in at least some situations. Secondly, it should be noted that there is a difference between some action's being rational in a given situation and its being rationally required. While it may be true that in any situation at least one action must be rational, it may not be true that only one (or even at least one) action must be rationally required. In some cases anything possible may be rational (not irrational) or any of a number of actions may be so.

Interestingly, Gauthier says:

My argument is then perfectly compatible with the view that it is rational to act egoistically whenever it is possible to act egoistically. There are no direct practical consequences to be drawn from the demonstration that it is not always possible to act egoistically (p. 455).

He nonetheless thinks that the incompleteness of egoism robs it of much of its appeal. If egoism needs to be modified, it loses

its simplicity and that simplicity, thinks Gauthier, is one source of egoism's appeal. But why suppose that the egoist would admit that his position needs to be modified? The egoist I have been considering is not committed to the view that every situation must be such as to allow him to carry out his egoistic calculations. He pursues his interests whenever he can and when he cannot he must simply hope for the best. There is nothing in Gauthier's argument to embarrass him, although there is a good deal to interest him.

(ii) Is egoism self-defeating?

Consider a Prisoner's Dilemma situation as follows:

| | Person P_1 does A_1 | Person P_1 does A_2 |
|-------------------------|-------------------------|-------------------------|
| Person P_2 does B_1 | (4, 4) | (8, -8) |
| Person P_2 does B_2 | (-8, 8) | (0, 0) |

The first number in the brackets represents the utility to P_1 of the outcome and the second number represents the utility to P_2 of the outcome. If each person reasons in a purely egoistic manner each, being fully acquainted with the situation, will perform that action which leaves him as well off as possible, given the action of the other. Thus P_1 will reason that if P_2 does B_1 he is best off doing A_2 and if P_2 does B_2 he is still best off doing A_2 . That is, regardless of what P_2 does P_1 is best off doing A_2 . But since P_2 can reason similarly to the conclusion that he is best off doing B_2 the outcome achieved will have utilities (0, 0).

Clearly, however, both would be better off if P_1 did A_1 and P_2 did B_1 since that outcome has utilities (4, 4). But how could egoists achieve this result? The answer is that as long as their actions are independent, neither's choice of action influencing

the other, they cannot achieve it. If each knows the other to be an egoist, each can anticipate the other's reasoning and $(0, 0)$ is the only possible outcome.

Suppose they consider agreeing to act together to achieve $(4, 4)$. As long as neither places any utility on keeping his agreements, no agreement is possible since each knows that no verbal exchange (agreement making) will make any difference. The redescription of action B_2 as " P_2 's breaking an agreement" will make no difference to the utility matrix. The conclusion is that there are utilities which are available to people who can trust each other to keep agreements which are not available to egoists (again assuming the egoists place no utility on the keeping of agreements).

This result might seem to provide an argument to the effect that there is a sense in which the egoist's policy of always trying to do as well for himself on every occasion as possible is a self-defeating one. By adopting this policy he cuts off the possibility of gaining the advantages of agreements in situations such as the one just considered. But some care is necessary in interpreting the result. First, we have assumed that the egoist does not place any utility on the keeping of his agreements and it is not clear that this possibility should be ruled out on Gauthier's definition of the egoist. (This possibility can be ruled out on my definition of egoism and hence I will not consider here how it might help or hinder the argument.)

Secondly, we have assumed that each of the agents in the Prisoner's Dilemma situation knows the other to be an egoist.

Our conclusion was, recall, that egoists, known to each other as such, cannot reap the advantages of agreements because neither can take the other's "agreement" seriously. Suppose, then, that we imagine the egoist's policy operating in a different situation; namely, one in which the other person in the Prisoner's Dilemma situation is known to the egoist not to be an egoist. In fact, suppose the egoist knows the other will keep his agreement as long as he believes the other parties to an agreement will keep faith. Then, finally, suppose the non-egoist believes the egoist to be like himself. Now the egoist can trick the non-egoist into agreeing that each will aim at the outcome (4, 4). The egoist (say P_1) will then perform action A_2 reasoning as before while the non-egoist (P_2) will perform action B_1 since he expects P_1 to keep the agreement. The resulting outcome has utilities (8, -8). Clearly then under some conditions there are utilities which may be open to the egoist which are not open to persons who can be trusted to keep agreements! There is no contradiction here with the earlier conclusion that there are utilities which may be open to persons who can be trusted to keep agreements which are not open to egoists. What utilities are open depends in part on the characters and beliefs of the other persons in a situation.

Gauthier concludes, correctly, in "Morality and Advantage":

[N]o man can ever gain if he is moral. Not only does he not gain by being moral if others are prudent [egoistic], but he does not gain by being moral if others are moral. For although he now receives the advantage of others' adherence to moral principles, he reaps the disadvantage of his own adherence. As long as his own

adherence to morality is independent of what others do, he must do better to be prudent (p. 469).

As long as others do not react to the egoist's egoism (as they will not if, for example, he deliberately deceives them or if they simply assume he will act morally) the egoist does not necessarily act in any self-defeating way by attempting to maximize his expected utilities on each occasion where this is possible.

Interestingly, Gauthier offers an argument designed to upset this conclusion in a later article, "Reason and Maximization."

There he argues:

The straightforward identification of rationality with the aim of individual utility-maximization [egoism], although not inconsistent, is nevertheless inadequate, because it denies the possibility of agreements which require one or more of the parties to refrain from the maximization of individual utility, yet secure to each of the parties greater utility than is possible without such agreement (p. 427).

Of course someone who identifies rationality with the aim of the egoist does not necessarily deny the *possibility* of such agreements but only the rationality of entering into them. Thus the sheer possibility of persons' entering those agreements does not show there is an inadequacy in egoism and Gauthier's argument does not suggest he thinks otherwise. What he attempts to show is that someone who identifies rationality with the aims of (straightforward, egoistic) utility maximization acts in a self-defeating way since there is another conception of rationality which opens the possibility of greater utilities to those who employ it. The allegedly superior conception of rationality is

one according to which it is rational to enter into agreements of the kind we found impossible for two egoists in the Prisoner's Dilemma. The crucial claim then is the following.

[S]ince the constrained maximizer [the adopter of the superior concept of rationality] has in some circumstances some probability of being able to enter into, and carry out, an agreement, whereas the straightforward maximizer has no such probability, the expected utility of the constrained maximizer is greater. Therefore straightforward maximization [egoism] is not self-supporting; it is not rational for economic man to choose to be a straightforward maximizer (p. 430).

I have not indicated exactly what is involved in the superior conception of rationality which endorses constrained (agreement-bound) maximization because I think the fault in Gauthier's argument is evident without going into the details of it. Indeed it is the same fault we noticed earlier in the attempt to use the egoists' "problem" in the Prisoner's Dilemma as an argument for the self-defeating nature of egoism.

It is true that in *some* situations the egoist (now understood as the person who identifies rationality with the aim of straightforward utility maximization) cannot reap the benefits of entering into agreements. But it is equally true that in some situations the egoist can reap the benefits of appearing to enter agreements and defaulting, and those benefits are not open to the person who adopts and acts on the conception of rationality which requires him to keep those agreements. Which of these sorts of situations are likely to occur more frequently and which conception of rationality is more likely to yield the highest utilities in the long run turns out to be an empirical

matter. It will depend on the conception of rationality others adopt, the devotion to rationality others display, the ability of the egoist to deceive others, and so on. Gauthier's argument would have considerably more appeal if the question of the adequacy of a conception of rationality turned, as he seems to think it does, on some necessary relationship among expected utilities rather than on such contingencies as those indicated above.

It should be noted that the egoist is not committed to advocating or advertising his conception of rationality. To require this would be like requiring that he announce his intention to break the agreement he wishes to seem to enter and that, of course, would trivialize the matter by making independent action impossible in the interesting cases.

My conclusion, then, is that there is nothing necessarily (either logically or empirically) self-defeating about the egoist's position as Gauthier understands it. Consequently there is no reason to think there is anything necessarily self-defeating in the egoist's strategy as I understand egoism.

8 Subjectivist egoism

So far I have considered only objectivist forms of egoism. The ethical egoist and the personal egoist both accept the idea that there are objectively valid standards for the assessment of actions, that some sorts of things really provide anyone with reasons for action regardless of their actual desires, wants, and so on. There is however, a kind of person we can call the sub-

jectivist egoist who does not adopt this objectivist posture. Vis-a-vis his status as an amoral agent he is in the company of other archetypical amoralists who share his subjectivism.

If the subjectivist egoist holds his position reflectively he may be expected to have some theory about the nature of practical reason. If he does not he may be described simply as someone who just does concern himself with his own interests without recognizing any need to suppose that everyone else does, or ought to, or has good reason to, concern themselves likewise. We may suppose that he considers himself rational but he may allow that other orientations are rational as well. That is, unlike the objectivist egoist, he does not think that anyone who is not an egoist is necessarily making any sort of mistake.

Like the objective personal egoist, the subjectivist egoist is open to the sort of mundane considerations advanced by Nielsen but because he is not disposed to demand that someone prove that he is actually *mistaken* about what his reasons really are, these considerations may have a greater impact and they may serve to effect a broadening of his interests. They may not. In any case it remains debatable whether simply expanding the concerns of a person to include direct interest in the welfare of others is sufficient to transform him into a moral agent. This is a question I will take up in a later chapter but for now let me just say that I think something more is required and that the something more has to do with a person's acceptance of these extra-regarding concerns as in some way required, valid, and so on. Simply coming to care about others, if this is felt as a matter

of personal taste, for example, does not transform the subjectivist egoist into a moral agent.

In many ways the subjectivist version of egoism is more interesting in the present context than the objectivist one. Unfortunately I have not said enough about subjectivism in general to develop this position very usefully or completely here. The notion of subjectivist egoism should become much clearer with the discussion in the last section of the next chapter and I beg leave to postpone further discussion until then.

9 The rationality of egoism

Personal and subjectivist egoism, in the light of the considerations advanced so far at least, seem to be possible and rational policies at least to the extent that they avoid practical and theoretical inconsistency. However, it might be argued that they are not rational if a broader view is taken of what rationality involves. In particular, the concepts of defensibility and reasonableness could be invoked.

As we have seen, the optimal strategy for the egoist involves his keeping his egoism secret to a large extent, at least insofar as he needs to rely on others to interact with him in certain ways. This means that he cannot, without acting against himself, publicly *defend* his policy or the actions which it dictates. To do so would involve revealing himself for what he is and that would be irrational given his concern to do as well as possible for himself. Thus, in one sense personal egoism is indefensible but in another it is not. The egoist can be

defended against charges of strict irrationality, but he cannot very well defend himself. It is difficult to argue with someone who cannot admit his own position to be as it is because one cannot find out *who* he is. The self-confessed egoist is already doing all that can be done against his position and the secretive egoist may not be vulnerable on any grounds he will consider.

If egoism is indefensible in a sense, it is also unreasonable *in a sense*, if to be reasonable is to be willing to enter into a sincere debate on the pros and cons of one's proposed actions. This involves a willingness to allow that there are publicly accepted and acceptable standards for what constitutes good reasons. The only reasons the egoist accepts are ones of self-interest and whether or not others would accept those reasons from him as justifying his actions or even as making them rational need be of no interest to him. Still, in another sense the egoist is reasonable. He acts for reasons and as long as others are prepared to give the right sort of advice he can respond with intelligent debate. Usually, however, when the charge of unreasonableness is made it is directed against someone who it is felt does not take seriously enough the effects of his actions on others - the departure he makes from the publicly accepted standards for reasons is in his refusal to take the interests of others sufficiently into account.

Another striking peculiarity of egoism under the formulation I have been using, according to which the egoist is concerned only with things like his own wealth, happiness, security, etc. is the exclusion of any interests which do not bear either

directly or indirectly on what can be called his welfare. Almost no one has concerns so narrowly confined. As has been noted often enough, it is scarcely possible to find a single person who takes no direct interest in the welfare of another, who would make no sacrifice of personal good for some cause of which he is not the beneficiary. In general, the fact that one feels satisfaction in performing actions which benefit others does not reduce such actions to pursuits of self-interest since it is the fact that one already cares about others which must be appealed to in order to explain the feeling of satisfaction. I will not deny the possibility of someone's being a true egoist, but it seems a remote possibility and, one could reasonably predict, an unstable one.

III

THE GROUNDS OF MORAL AGENCY

1 The possibility of amorality - non-rational capacities

In the last chapter I examined egoism as one kind of amorality. I argued that, while ethical egoism is open to serious attack, especially as regards its practical consistency, there are two versions of egoism - (objectivist) personal egoism and subjectivist egoism - which can survive the arguments usually directed against egoistic theories. There are, however, other arguments which purport to show that moral agency is, in some sense, a requirement on all mature, rational, fully human persons. In this chapter I will be dealing with a number of these arguments. Some of them are directed against objectivist amorality and are designed to show that one is bound (somehow) to take account of the welfare of other persons in one's deliberations about action. Others are directed against subjectivist amorality and these are calculated to demonstrate that anyone who fails to recognize objectively valid standards of action or objectively valuable states of affairs is making some important sort of mistake.

Consider the question, "Why be moral?" In the introduction I suggested that there are three senses of "moral" depending on

whether "moral" is being contrasted with "non-moral," "amoral," or "immoral." Now obviously the sort of answers which might be given to the question, "Why be moral?" will depend on how the word "moral" is understood. If the question is "Why be not non-moral?" it makes little sense because the non-moral agent (taking "agent" broadly) will lack the capacity to consider the question intelligently. Animals, the severely retarded, and the thoroughly insane, for example, could not face the question of whether to become other than they are.

If the question is understood to be "Why not be an immoral person?" then, or so I have argued, it can arise only for someone who already is a moral (not amoral) agent - i.e. for someone who, at some level, already recognizes that there are moral requirements. The moral agent might ask the question, "Why not be 'immoral'?" where he wants to know what reasons can be given for acting in accordance with what are *taken*, generally or in his social group, to be the requirements of morality. And he can request to be *motivated* to do what he already admits to be required morally, but he cannot in any other way raise the question seriously, unless, of course, he is actually considering abandoning morality and taking up an amoralist outlook. But is amorality, after all, a real possibility?

There are two main directions from which the impossibility of amorality can be argued. First, it might be thought that being fully human (where I mean by "fully human" just not non-moral) involves, among other things, having a certain sort of psychological sensitivity which necessarily manifests itself in the use of

moral language and in a preparedness to think and speak "in the moral mode." Moral sense theories, for example, suggest this approach. Secondly, it has been argued that being rational, traditionally a defining characteristic of human beings, involves a person in morality, at least insofar as he follows his rational nature. The Kantian approach comes readily to mind. It is difficult to deal with these two lines of thought without considering the details of the particular conceptions of moral sensitivity and of rationality held by actual theorists; however a brief look at some well-known attempts to deal with the question may be of some use.

David Hume believed that the ability to make, and the willingness to use, moral distinctions is universal among human beings in virtue of their natural emotional make-up. He says:

If any man from a cold insensibility, or narrow selfishness of temper, is unaffected with the images of human unhappiness or misery, he must be equally indifferent to the images of vice and virtue: As on the other hand, it is always found, that a warm concern for the interests of our species is attended with a delicate feeling of all moral distinctions; a strong resentment of injury done to men; a lively approbation of their welfare. In this particular, though great superiority is observable of one man above another; yet none are so entirely indifferent to the interest of their fellow-creatures, as to perceive no distinctions of moral good and evil, in consequence of the different tendencies of actions and principles.

(Hume, 1966, *Enquiry Concerning the Principles of Morals*, p. 60; emphasis mine)

If one is prepared to grant Hume's claim that "perceiving moral distinctions" is a matter of preferring (feeling approbation of) human happiness and prosperity over unhappiness and suffering,

then the fact, if it be such, that everyone does, other things being equal, prefer happiness to unhappiness, whether in themselves or in others, will indeed make problematic the claim that amorality is a possibility for fully human persons. On Hume's view, to be amoral would be to be completely indifferent to both human happiness and human suffering in general. Since there are few people, if any, in this category (save perhaps some who are arguably deficient in some of the capacities outlined earlier as being involved in someone's being fully human), the category "amoral agent" must be sparsely populated indeed. But, as I tried to show in Chapter I, there is much more to morality than Hume and emotivists in general can easily account for. At the very least, the "feeling of approbation" must be a very special sort of thing and must be understood along the lines of Roger Scruton's analysis of a moral attitude (see Chapter IV). Thus it is not possible to show that the amoralist will necessarily be unconcerned with human welfare either in general or in the case of particular persons, although it may be that the converse is true. That is, it may be that anyone who is unconcerned with anyone else's welfare is amoral if he is not actually non-moral.

Butler, like Hume, thought that some concern for the well-being of others is virtually universal among human beings and constitutes an essential aspect of human nature. Yet Butler realized that something more is required than a feeling of approval in order to account for the peculiar authority of moral judgments, even if those judgments are based on reactions shared

by all human beings. There is:

... a superior principle of reflection or conscience in every man, which distinguishes between the internal principles of his heart, as well as his external actions: which passes judgment upon himself and them; pronounces determinately some actions to be in themselves just, right, good; others to be in themselves evil, wrong, unjust

(Butler, 1967, "Upon Human Nature,"

II, *Fifteen Sermons*, p. 53)

Conscience is a faculty of fully human persons and it is absent in animals and perhaps also in persons not fully human through immaturity or some defect in their nature.

Unfortunately, Butler's view of the nature of conscience is less clear than we might wish. He conceives of it as operating to issue moral judgments, which stand in a relation of superiority to the promptings of the particular passions such as hunger, fear, love of another, and even to the two principles of benevolence and self-love. In a way, conscience introduces a new normative element into the deliberative process, but since Butler thinks that all passions properly understood promote both the good of self and of others, it is not obvious what the role of conscience amounts to other than to perform the apparently non-normative function of indicating wherein lies the way to achieve these allegedly coincident ends. He says at one point:

The sum is, men have various appetites, passions, and particular affections, quite distinct both from self-love and from benevolence: all of these have a tendency to promote both public and private good, and may be considered as respecting others and ourselves equally and in common

(Butler, 1967, "Upon Human Nature," I

Fifteen Sermons, p. 38)

If we are skeptical about Butler's claim that all our inclinations

do tend to promote both public and private good (and we are certainly entitled to maintain some reservations on this point), then the peculiar authority of conscience must derive from something other than its function of indicating which actions serve this dual purpose. In any case, Butler has not shown that his model of the deliberative process is even acceptable, let alone the only one possible for human beings. In Chapter IV, I will have more to say about deliberation, especially as concerns the possibility of avoiding the moral categories in which conscience supposedly issues its decisions.

There are, of course, many other theories according to which being fully human involves having certain psychological sensitivities whose function it is to register distinctions which are essentially moral in their import. All such theories, I think, fail to solve the problem of showing that the "facts" which these senses or faculties register are necessarily loaded morally in the required way. Concepts like "obligation," "duty," "right," and "ought" are not easily teased out of emotions, sentiments, and the like. Intuitionist theories like those of H.A. Prichard (1912) and G.E. Moore (1903), which take a more direct route to moral distinctions than moral sense theories, are little better off. Their difficulties stem from the fact that it is difficult to make such direct appeals very satisfying intellectually and the fact that so many apparently otherwise capable people are unable to locate in their experiences anything like intuitive insight into moral truth.

2 The possibility of amorality - conceptual structures and rationality

If it is not in virtue of some inevitable morally significant psychological sensitivity that it can be shown that amorality is in some sense impossible, perhaps the thesis can be argued from an analysis of the nature of rationality, or some conceptual consequence of someone's having the capacities which place him in the realm of the fully human. Using the term "rational" in an extended sense to cover both these areas, we can ask if every rational agent is necessarily a moral agent. Is there something conceptually amiss with someone who, while perfectly capable of moral thinking, nonetheless does not apply moral concepts?

I suggested earlier that the objectivist amoralist (the Nazi fanatic, the religious zealot, *et al.*) could speak more or less comfortably about what people ought to do, about their obligations and duties and about what sorts of actions are right and wrong. (The egoist has problems here because he cannot freely promulgate his views unless he holds the mistaken belief that he will attain his own ends better in a world of egoists.) These concepts bespeak an acceptance of an objectivist posture, and, while they are not properly moral concepts in the mouth of an amoralist, he (and we) can understand them well enough. As I hope will become clearer as I proceed, the subjectivist amoralist cannot use such moral or quasi-moral concepts very easily or sincerely. He may have those concepts in the sense of being able to understand, in large measure, what others intend

to convey with them, but he has no sincere use for them to express anything he wants or needs to say.

In general, it is possible for someone to have a concept or set of concepts, in the sense of being *capable* of using those concepts in thinking about his experiences, actions, etc., without his actually *using* the concepts. An atheist may, for example, understand religious concepts such as "grace," "sin," "redemption," "God," and so on, and yet not use those concepts himself (except perhaps in the sense that he uses them to understand what religious persons say when they use them). To fail to use a concept in this way is different from failing to use concepts like "unicorn" and "tooth fairy." Rather, it is more like not using concepts like the existentialist concept of "authenticity" or the concept of "existence outside time." I am not sure I can clarify the difference I seem to detect among ways of not using concepts, except to say that in the sense in which I am claiming an amoral agent does not use moral concepts, his not doing so consists largely in there being no room in his way of thinking for him to operate with moral notions. He cannot see himself or anyone else as *being under a moral obligation*, for example, roughly in the way an atheist cannot see someone as *having sinned* since there is no way for this to fit into his way of thinking about the world and his place in it. Being an atheist is not, I take it, simply a matter of disbelieving a proposition about the existence of a certain entity and that is why there is more to a religious conversion than simply coming to believe a proposition. A religious convert comes to think in

a different way; he does not just think different things are true.

Different cultures, different sub-cultures and even different individuals have, to some extent, different conceptual structures. It is not a simple matter to say, in general, what constitutes differences in conceptual structures. In some cases, the variance consists in little more than one person's or group's having, while another lacks, a concept or set of concepts, but sometimes there is much more involved. Concepts change, become more complex, and exhibit altered logical and quasi-logical interrelationships. To say just that much is not very helpful or illuminating and to say much more is very difficult. The point at which this bears on the present discussion is that what it is rational for someone to think, believe, and do is in some measure dependent on his conceptual structure. That is not to say that there is never anything to choose, on grounds of rationality, among various conceptual schemes, but only that what \ast is a rational procedure for one person may not even be a possibility for another.

Suppose, for example, that there were a society in which people conceived of themselves not as individuals in the way we do, but rather as essentially a part of a social "we," even to the extent that they were incapable of formulating a clear idea of themselves as distinct, unique, independently existing individuals. Such persons may have no way of formulating the notion of amorality, of self-interest or even of morality. They would be bound by their conceptual apparatus to thinking in a way we might call "moral," without their realizing it, having, as it were, no way of contrasting their mode of thought with

an appropriate alternative. For such persons, to ask "What shall I do?" may be to ask "What is required of this part of us?" or something like that. Having no way of considering the idea that what is socially valued or valuable need not coincide with what an individual values, any attempt to even formulate the idea of an alternative to morality would likely fall into apparent absurdity.

Now the subjectivist amoralist as I conceive of him is not afflicted by the kind of conceptual isolation illustrated in this example. He does understand moral and quasi-moral concepts but he does not use them. However, it is not just concepts such as "morally ought," "obligation," "right," etc. which are relevant here. If it is to be argued that the subjectivist amoralist is irrational we must also attend to the concept of rationality itself and perhaps to others as well.

In evaluating a person's rationality we attend to his use of reasons. There are two more or less distinct aspects of our evaluations; we can assess the constancy and consistency of his reasons and we can assess the contents of his reasons themselves. The former assessment is primarily a matter of consistency among a person's reasons (and beliefs) and the latter of the intelligibility or acceptability of his reasons themselves. We may say that someone who takes fact A to be a reason for ϕ -ing but does not take fact B which differs from A only trivially and who cannot show that there is any relevant differences between the facts, is irrational because inconsistent. On the other hand, we also say that someone who treats facts of an utterly (apparently) irrelevant sort as a reason for acting (or believing) is irrational.

In both cases we fail to understand a person's actions or beliefs adequately but in the former case (involving logical contradiction or inconsistency) there are strong grounds for saying that our failure to understand is attributable to the inherent incomprehensibility of the phenomena, whereas in the latter case it is more plausible, at least sometimes, to attribute the problem to an incongruity between the conceptual outlook or value structure of the person trying to understand and the person whose behaviour seems irrational. I will be discussing a subjectivist view of reasons further in Chapter V and the relevance of these brief comments should become clearer then.

Beginning in the next section I will be looking at various attempts to show that anyone who is not a moral agent is irrational, but first I want to consider a recent controversy in which the issue is whether or not it even makes sense to ask for reasons for being moral (as opposed to being amoral).

3 Does "Why be moral?" make sense?

Often the senses of "moral" I have tried to keep separate are not very carefully distinguished and so treatments of the above question sometimes run at cross purposes.

Stephen Toulmin has argued that the question "Why be moral?", if it is understood as a request for a justification for doing some particular action which is already allowed to be in accordance with a current social practice (and on Toulmin's view that is sufficient for an action's being morally right or required), or if it is understood as a general request for a reason for doing what one sees to be morally right, has no literal meaning.

The question, if not absurd, is a "limiting question," that is, one which seems to require an answer of a kind which cannot be given. He writes:

When it has been pointed out that an action conforms unambiguously to a recognized social practice, there is no more room for the justification of the action through ethical reasoning ... (p. 217).

When someone asks, perfectly generally, 'Why ought one do what is right?', and is not satisfied with the answer that the sentence, 'You ought to do what is right', expresses a truism, his question is also a 'limiting' one.

(Toulmin, 1968, *Reason in Ethics*, p. 218)

Limiting questions of this type, claims Toulmin, must be understood, not literally, but rather as a request to be given some motivation for doing the act or type of act in question.

Of course, people do sometimes fail to feel motivated to perform actions which they accept as morally required, either because they also have non-moral reasons for doing otherwise or through sheer laziness. Someone could use the question, "Why be moral (in this case)?" to request motivation, but Toulmin does not take very seriously the possibility that someone might question the whole enterprise of thinking morally. This is particularly odd in view of Toulmin's account of what's involved in moral justification. If to say that something is morally required or justified or right is essentially to claim just that it accords with some social practice, it seems quite possible for someone to question morality on some general grounds (eg. of rationality). "Why be moral?" will seem like a very strange question in the mouth of a moral agent who is using "moral" sincerely in its "not

immoral" sense, because in granting that some action is morally required he has already accepted all of the justification necessary for doing it.

Kai Nielsen (1968), on the other hand, takes the possibility of an amoral perspective seriously and argues against Toulmin on the grounds that there are many kinds of practical reasoning besides moral thinking. It is not self-evident, Nielsen points out, that anyone must give moral reasons priority or supremacy over other sorts of reasons. He construes the question "Why be moral?" as the moral skeptic's question "Why not be amoral?" This question is obviously not a moral one.

In a reply to Nielsen's article, R. Beehler sets out a second sort of reason for rejecting the question as being a mistaken one. He argues that one cannot *decide* whether or not to become or remain a moral agent.

How is man to adopt the 'moral point of view'? How is he to *decide* to be kind, unselfish, courageous, honest, loyal, just, and so on? For that is what 'having the moral point of view' is. The moral point of view is not some point above honesty, selflessness, mercy, integrity, justice, in terms of which those are seen to be worthwhile. Having the moral point of view is just approaching life and persons honestly, justly, forgivingly, and so on, where these things matter to you, where, in a sense, these things are what you are. But if you don't have a regard for, say, honesty now, how are you to decide to have one?

(Beehler, 1972-3, "Reasons For Being Moral," p. 16)

There is a very important point here which never comes explicitly into the debate between Nielsen and Beehler. It is Beehler's assumption that being a moral agent is entirely a matter of

having certain virtues, character traits, or inclinations, values, etc. and that this exhausts the notion of the moral point of view. This is the sort of claim I objected to earlier in Hume's analysis of morality. I think that Beehler is wrong in supposing that being a moral agent is simply a matter of certain things mattering to one. Of course, certain things will matter to a person who is a moral agent to the extent that he takes seriously the results of his moral thinking, but there is no obvious reason why at least some of those things could not also matter to someone who did not appreciate things morally. There may be some virtues which are specifically moral in the sense that they involve values which only moral agents could have since they involve the use of specifically moral concepts. Justice, under some analyses, may be one such. Furthermore, it is not easy to see how "having the moral point of view" can be analyzed adequately in terms of what matters to one unless some deeper analysis is available of why just these things count as moral concerns. But suppose that there is something which underlies these moral concerns - perhaps some concern such as for the harmony and welfare of mankind - it still remains to show that one cannot decide, in an appropriate sense, to care about such things.

Probably no one will suggest that we can decide to care about things, in the same way we can decide to go to a movie, but surely nothing that strong is required in order that the question of whether or not to become (or, for most people, whether or not to remain) a moral agent, make sense. All that is necessary

is that there be the possibility of having reasons one way or the other and of those reasons being instrumental in producing, perhaps over time, the existence or extinction of the attitude in question. Deciding to be moral is more like deciding to be scientific than like deciding to believe specific scientific propositions. One can give up mysticism in favour of science for reasons in much the same way that one can give up morality in favour of self-interest. Thus, even where it is strained to speak of deciding to do this or that, it may nonetheless make sense to raise the question of whether there are reasons for taking a modified view of the matter, perhaps with a view to effecting a gradual change in one's perceptions, emotions, and beliefs. Indeed, Beehler finally came to admit:

If deciding not to be moral is resolving to try not to be moral, then a man could, I suppose, do this. He could decide to try not to appreciate things morally.

(Beehler, 1972-73, "Morals and Reasons," p. 20)

Still, I think, ceasing to be a moral agent may involve more than, or rather something different than, ceasing to care about certain things. It involves ceasing to care about things morally, ceasing to use moral language and ceasing to engage in typically moral activities. Furthermore, an amoral outlook may be a consequence of someone's becoming convinced that it is a rational alternative to moral agency and thus may not require any special effort.

Because moral judgments look like ordinary statements it is easy to fall into thinking that one must either believe them

true or believe them false or else withhold belief and disbelief while allowing they must be one or the other. There is, however, another possible attitude. Someone may not have anything to say which needs to be expressed in moral judgments. In a sense, this is like deciding that no one has any moral obligations but to put it that way is apt to make the position seem too close to that of someone who decides that (believes that) no one has a moral obligation to submit to the military draft, where the latter is a moral position.

Finally, against the intelligibility of the question "Why be moral?" it has been argued that the only ultimate answers to practical questions are ones which derive from either moral considerations or considerations of self-interest, and to offer reasons of the first sort is to fail to give an answer at all, the attempt being circular, and to offer reasons of the second sort is to attempt the impossible since a part of the *raison d'etre* of morality is that it provides practical considerations which take precedence over those of self-interest. The only time that the question of whether or not to be moral will arise, on this view, is when morality is in conflict with precisely those things it is supposed to override. This sort of argument, however, is applicable only if the question is "Why be moral and not immoral?" i.e. only if it is assumed that the question concerns the alternatives *moral* vs. *immoral*. This question cannot sensibly arise for a moral agent except insofar as he is considering the possibility of amorality, but in that case the fact that it is part of the *raison d'etre* of morality that it provides

overriding reasons for actions will be quite beside the point. The absurd question then, is not "Why be a moral agent?" but rather the question, "Why should I perform this action which I see to be morally required?" as asked by someone who is a *dedicated moral agent*. As Toulmin argues, the only sense in which that question, asked by a dedicated moral agent, can be understood is as a request to be motivated to do that which it has already been allowed is the only justifiable thing to do. Of course, the latter question cannot arise for the amoral agent since its statement requires the sincere use of concepts which he does not use. To judge an action to be morally right is, among other things, to accept a justification for doing it and to be motivated, other things equal, to do it.

Hopefully I have vindicated the question with which this chapter is concerned with from the charge of absurdity. I do not claim, however, that this is sufficient to show that it necessarily has a straightforward or compelling answer one way or the other. I now turn to an examination of some attempts to show that there are reasons for being a moral agent which derive from rationality or other deep-seated features of our conceptual apparatus.

4 The Golden Rule

The idea that any rational agent must be open to arguments of the Golden Rule sort is a powerful one. The Golden Rule expresses a central aspect of moral thinking and the claim that it is also linked closely with the concept of practical ration-

ality is the basic contention in Kant's treatment of the Categorical Imperative (1964). Perhaps it would be better to say that Kant's Categorical Imperative represents a minimal interpretation of the Golden Rule. It is generally agreed that there are very serious problems involved in saying exactly which version or versions of the Categorical Imperative Kant is entitled to and how much of morality is captured by them. Certainly there is something to the notion that insofar as an action is performed deliberately it can be said to embody some maxim or principle and that it would be irrational for someone to act on a maxim which it would be impossible for any number of other persons to act on. A maxim which is rational (not irrational) for one must be rational for all and it must be possible for everyone to act rationally - at least I am not prepared to argue otherwise.

Kant's formulation of the Categorical Imperative, of course, is stronger than this. It requires not merely that one act according to maxims which could be acted upon universally, but that one act only on maxims *that one could will* that all should act on. This takes us closer to the Golden Rule since it corresponds to the rule (the Silver Rule we might say): "Do unto others as you would will they do unto you." The Golden Rule, however, requires one to consider what maxims one would have others act upon (at least insofar as their doing so has an effect on one) and then to act oneself on those maxims where they apply. It requires that we act only according to maxims which we would have others act upon as well. It is clearly a moral rule in

that it operates to direct attention to the effects of one's actions on others and bids us choose maxims on the basis of how we would react if the roles were reversed, so that we not only see ourselves as affecting others but also as being affected by our maxims and the actions they sanction.

It is apparent, if my attempt to sketch some of the main features of morality was even reasonably successful, that the Golden Rule, but not the Silver Rule, can plausibly be considered to express the essence of moral thinking. For example, the requirement that one be prepared to defend and promulgate one's moral beliefs makes sense if one is acting according to a rule which requires that one will that everyone act on the same maxims one uses oneself. The question to be addressed here is whether or not there is anything irrational or otherwise conceptually untoward in someone's failing to choose his practical principles in the way suggested by the Golden Rule. In order to discuss this question entirely satisfactorily, I suppose I would have to present and defend a conception of rationality, and I do not pretend to have a complete theory of rationality by any means. However, it will be agreed that being rational involves, at the very least, avoiding certain sorts of inconsistency. In the practical realm, a rational agent does not knowingly perform actions which are self-defeating (inconsistent with the ends for which he undertook them) and he does not perform actions which are inconsistent with his practical judgments. For example, he does not do A when he judges that he has (conclusive) reason not to do A.

One fairly straightforward argument for the rational required-ness of following the Golden Rule can be disposed of right away. It might be argued that, in situations "calling for a moral resolution," to do something deliberately in itself commits one to thinking that one is morally justified in doing it. But, the argument goes, that commits one to the view that anyone else would also be acting in a morally permissible manner in acting similarly (in adopting the same practical maxim). Since it is in the nature of a judgment of moral permissibility that one must will that others, as well as oneself, act in a morally permissible manner, one must will that the maxim of one's action be universally adopted. The problem with that argument, even if we ignore the implicit claim that permissible action in any situation can result only from using a unique maxim, is that it assumes that one must already be committed to evaluating one's own and everyone else's actions morally and that is just the point in dispute.

Perhaps a similar argument could be constructed in morally neutral terms, however. Indeed, I think C.I. Lewis has attempted just this. First, it should be noted that Lewis does not see the difference between the Golden Rule and Kant's Categorical Imperative which I suggested above and which emerges from a concentration on the possibility that someone may be able to will what he does not in fact and never would will. That is, I take Kant to have argued for, and only to have been entitled to, the claim that purely in virtue of being rational one is constrained not to act on maxims one *cannot* as a rational agent will.

to become everyone's maxim. I realize that there is another way of understanding Kant on this point which rests on the idea that if one can will something (can work oneself into the frame of mind where one would will it) then in some sense, one does will it. It makes a difference if the "can" is a logical or a psychological one; in the latter case it is not so clear that there is a great difference, or at any rate a relevant one, between what one can will and what one does will. It is as though to come to see that one can will something is to have undergone all that is necessary to have come to will it. In any case, I think that the logical "can" is all Kant is entitled to since he is concerned with rational beings, some of which may have no psychology as we understand it.

Lewis (1969, p. 75) formulates the Golden Rule thus: Do no act which contravenes any rule which you would call upon other men to respect and conform to." His argument for the rational requirement to accept this rule is advertised as a *reductio ad absurdum* of the repudiation of it. What, then, is involved in repudiating it?

What the Categorical Imperative says is no more than that there is a non-repudiable distinction between the right and wrong, which affects whatever we must decide by deliberation. And this distinction is non-repudiable for any creature who thinks, and thinks to a purpose, and is called upon to decide by thinking. There is the moral imperative because, important amongst things the rational animal must decide, there is the question of how he shall behave toward his fellows. And any creature who talks to himself about that will find himself in the predicament of pragmatic contradiction if he says to himself that there is a way of acting which is right for him but wrong, in

the same premises of action, for another creature who likewise decides by talking to himself.

(Lewis, 1969, *Values and Imperatives*, p. 200)

Lewis' basic idea is that there is a point to deliberation; that in deliberation there is a right answer and a wrong answer, that some actions are justified and others are not, that there is something one ought to do and something one ought not to do. Furthermore, since the criteria for everyone are the same, it being absurd to suppose that two persons "in the same premises of action" might equally rationally apply different maxims or reasons to two different justified results, a rational agent is bound to allow that he ought to act only as he is prepared to see everyone act. To repudiate the Golden Rule, then, is to suppose that there are no criteria for choosing among possible actions in deliberation: it is to suppose that deliberation is pointless.

Let us grant that there is a point to deliberation in the sense that some possible actions for a given person can be said to be ones which he has reasons to perform and others to be ones which he has reasons not to perform. Clearly anyone who is similar to this person will, in similar circumstances, have similar reasons. In this sense, what is right for one is right for the other and it would be irrational to suppose that some person P has reasons which make some action A the one which solves his practical problem and that another similar person Q, similarly situated, has reasons which make some different action B the one which solves his practical problem. But Lewis must show more

than this. He must show that there is an irrationality involved in someone's acting on reasons he would not *call upon others* to act on. Certainly it would be dishonest to actually call on others (in the name of rationality) to act against their reasons, but one could remain silent, and anyway dishonesty is not obviously irrational.

The reason Lewis' argument fails is, no doubt, apparent by now. There is an important difference between (a) granting that the deliberations of others are governed by the same rules of rationality as one's own, and (b) taking some interest in whether or not others deliberate according to those rules. This difference corresponds to the difference between the Silver Rule and the Golden Rule. It turns out that Lewis' argument for the Golden Rule does not support it in the form he gives it above.

I do not see any way of showing that there is any practical or theoretical inconsistency in someone's allowing that there are general principles which describe the point of deliberation and which apply to all rational agents, and at the same time remaining outside the project of "calling on others" to use those principles. There does not seem to be any way of showing that it is in any way a part of the point of deliberating to attempt to find actions which embody maxims one hopes others will use. Since the general principles of rationality may make reference to the value structure of the agent, their implications for different agents in similar circumstances may be very different. That is, just what it is rational for a person to do may depend on what sort of person he is, but more of this later.

There is another strand in Lewis' attempt to establish the Golden Rule which should be considered. He says:

The basic imperative is ... simply that of governing oneself by the advice of cognition, in contravention, if need be, to impulses and the inclinations of feeling. And this imperative can be avoided only by the incapacity to deliberate and make decisions. This most comprehensive imperative of rationality may be called the Law of Objectivity

(Lewis, 1969, p. 167)

The Law of Objectivity is relevant to moral deliberation in requiring that one take full cognizance of the reality of other persons.

Granted this law of respect for objective fact as such, the ground of our obligation to another person becomes obvious, does it not? The reason for it is that we know him to be as real as we are, and his joys and sorrows to have the same quality as our own.... The principle of objectivity dictates compassionate regard for others just as, so to say, prudence dictates compassionate regard for my self of tomorrow.

(Lewis, 1969, p. 141)

If we recognize, as Lewis does, that the notion of "compassionate regard" here does not necessarily involve one in feeling compassion, as that is normally understood, but just the full recognition of the effects of one's actions on others, it is apparent that the ground of obligation provided by the Law of Objectivity is not as firm as is required by our ordinary conception of moral obligation (which does not allow, I think, that our obligations depend on our feeling compassion for others - although our being motivated to perform our obligations may be thus contingent).

Thus another defense of the Golden Rule may be reconstructed as follows. Rationality requires that, if deliberation is to have a point at all, we take into account in asking ourselves what to do, the full reality of others. If this compassionate regard for others operates in a certain way (roughly so that we come to feel the effects of our actions on others as poignantly as we feel their effects on ourselves), then we will, in fact, be led to act only on maxims which not only could be, and could be willed to be, acted on universally, but which we imagine others actually acting on (through the imaginative exercise of placing ourselves in their shoes and feeling the effects of being subject to those maxims). Thus, to act on a maxim rationally is, for someone who is sufficiently moved in a certain way by imagining the effects of his actions on others, to act in a way he wills to be a universally accepted way of acting. Clearly, then, it is a matter of psychology and not rationality whether the Golden Rule will get a grip on the deliberations of any particular individual. That is not to say that it is necessarily just a peripheral fact about human beings that the Golden Rule is capable of influencing action; it may not be. It would require a separate argument to show that someone cannot remain outside the jurisdiction of the Golden Rule and be a truly human being.

At the beginning of this chapter I said that what it is rational for a person to do may depend on his conceptual structure - on what concepts he has - and according to my reconstruction of Lewis it also depends on his psychological sensitivities. If

it could be shown that someone who does not take account of the effects of his actions on others as though they were effects on him does not have a conception of the equal reality of other persons, this would give some sense to the claim that the amoralist is cognitively deficient, if not actually irrational. I do not think this can be shown any more than it can be shown that in order to be said to have an adequate conception of the equal reality of animals a person must take their experiences into account, in acting, as if they were his own. But even if it could be shown, there remains the problem of showing how the distinctively moral notion of obligation, duty, and so on can be grounded in this requirement.

5 Thomas Nagel

It seems, then, that there is no way to proceed from the recognition of the equal reality of other persons to the conclusion that one must take their interests into account when one acts. At least there is no way of showing that a vivid perception of the effects of one's actions on others will necessarily result in a typically moral concern for others. At this point another approach suggests itself. Suppose we look at things a little differently and, instead of concentrating on the consequences of recognizing that others are real in the same way we are, we look at the consequences of having a concept of ourselves as just one person among others equally real. Actually this involves more than just looking at the same thing from a different angle; it is more like looking at oneself in a particular way in light

of a recognition of the equal reality of all persons. Thomas Nagel (1970) has attempted to exploit this possibility in his book *The Possibility of Altruism*.

Nagel's argument is complex and it is not possible to deal adequately here with everything in this fascinating book. In particular I shall ignore his effort to develop an analogy between moral reasoning and prudential reasoning according to which the former requires a conception of oneself as one person among other persons equally real while the latter requires a conception of the present as one time equally real with other times (especially the future). There are enough problems in the analogy to make its usefulness doubtful, in any case.

Nagel ranges his arguments primarily against what he calls subjective egoism, a view which is entailed by the position I am calling amoralism. He says:

The most philosophically attractive view (and hence the one most worth rejecting) is one which denies that reasons depend on the assignment of objective value to anything ... (p. 96).

Just what is involved in "the assignment of objective value" to something is not entirely clear. Nagel certainly is not committed to the Moorean view that there is a non-natural property of goodness such that just those things which have that property have objective value, although such a view is consistent with much of what he says. At the very least, to assign objective value to something, say some state of affairs, is to suppose that anyone has a reason to bring about or maintain that state of affairs. The subjective egoist (the amoralist) can allow that there may be, as a contingent fact, some universally valued state

of affairs, but if there is it will be in virtue of facts about individual persons primarily, and not just because of some features of that state of affairs in itself. Certainly not all things valued by someone will be valued universally and even where something is valued by everyone its value consists in just this contingent fact, so that it is never possible for an amoralist to "assign" objective value to something if this is to involve anything more than valuing it himself and/or believing that others as a matter of fact also value it, and it is clear that for Nagel more than this is involved, as we shall see.

The notion of objective value can be clarified somewhat by looking at Nagel's treatment of reasons for action. A reason is identified with a predicate which applies to actions. Thus, predicates such as "promotes the agent's welfare," "is an act of benevolence," and "conduces to someone's welfare" identify reasons or putative reasons for performing the actions to which they apply. Nagel places a universality condition on reasons which requires that reasons be the same, in some sense, for everyone. If one thinks that the fact that some predicate R applies to some action A gives some person p a reason to do A, then one must also allow that any other person to whose possible action B, R applies, also has a reason to do B. Another way of putting this point is to say that supposing that R's applicability to A gives p a reason to do A commits one to the practical principle:

(p,A) (If R holds of A, then p has a reason to do A)

i.e. "For all persons p and all actions A open to p, if the

predicate R is true of A, then p has a reason to do A." Using the phrase "*prima facie*" to indicate that the reason need not be a conclusive one, and allowing A to range over states of affairs, events, etc., as well as actions, we get:

(p,A) (If R holds A, then p has *prima facie* reason to promote A).

Now, according to Nagel, reasons and practical principles can be either objective or subjective according to whether or not a reference to p, the agent, occurs in the formal statement of the reason predicate or practical principle. If there is a reference to the agent, the reason or principle is subjective; otherwise it is objective. The left hand column below contains examples of objective reasons and principles and the right hand column examples of subjective ones.

- | | |
|---|---|
| 1. that A will make Canada great | 1'. that A will make p's country great |
| 2. that A serves G.D.'s interests | 2'. that A serves p's interests |
| 3. that A makes someone happy | 3'. that A makes p happy |
| 4. Everyone has a reason to promote Canada's greatness. | 4'. Everyone has a reason to promote his country's greatness. |
| 5. Everyone has a reason to promote G.D.'s interests. | 5'. Everyone has a reason to promote his own interests. |
| 6. Everyone has a reason to make people (himself included) happy. | 6'. Everyone has a reason to make himself happy. |

Since objective principles may entail subjective ones (eg. 6→6'), it will be useful to stipulate that when I refer, for example, to persons who accept subjective principles, I will mean that they accept *only* subjective principles. Indeed, let me specify that

any reference to a subjective principle or reason is to be understood to concern a reason or principle which does not rest on an objective one; thus, whether it is objective or subjective will depend not only on its form, but also on the practical thinking of the person(s) supposed to accepted it.

A *practical* judgment, for Nagel, is a judgment like "p has a reason to \emptyset ," or "p has a reason to promote A," and so on. Nagel includes judgments of the form "p ought to do A," but I prefer to deal in the former sort since I think Nagel does not pay sufficient attention to the fact that these two sorts of judgments sometimes have different meanings in significant respects. A very important concept for Nagel is that of a practical judgment's having *motivational content*. He says that "the acceptance of such a judgment [one with motivational content] is by itself sufficient to *explain* action or desire in accordance with it, although it is also compatible with the non-occurrence of such action or desire" (p. 109). First-person practical judgments, such as "I have a reason to go to the store," have motivational content because they are "not merely classificatory: they are judgments about what to do; they have practical consequences" (p. 109). A third person judgment such as "John has a reason to stay in bed," then, has motivational content if the person who makes it (accepts it) thereby opens the way for an explanation of his acting to promote John's staying in bed or perhaps only his desire that John do so.

Nagel states his project thus:

The thesis which I propose to defend is simply that the only acceptable reasons are

objective ones; even if one operates successfully with a subjective principle, one must be able to back it up with an objective principle yielding those same reasons as well as (presumably) others. Whenever one acts for a reason, I maintain, it must be *possible* to regard oneself as acting for an objective reason, and promoting an objectively valuable end (pp. 96-7).

I will turn shortly to Nagel's attempt to establish this thesis and to the connection between it and the notion of motivational content but first I should like to offer some brief remarks on the thesis itself. The first thing which strikes one is the oddity of the claim considered as a general constraint on *all* reasons. It seems quite incredible that anyone should suggest that any time one does something for a reason one *must* be able to regard oneself as promoting an objectively valuable end - one which *anyone* has a reason to promote. Note that what is required is not simply that there be some similar end which each person has a reason to promote (that is required, but it has to do with the universality of reasons and not with their objectivity). That is, the objectivity thesis is not that, for example, if I have a reason to wash behind my ears, then you do too - or at least you have a reason to do whatever in your case corresponds to my washing behind my ears, if my reason for doing so is derivative in the sense of falling under a more fundamental reason. The objectivity thesis requires that if anyone allows that I have a reason to wash behind my ears then everyone has a reason to wash behind my ears, or perhaps to see that I do so, depending on the form of the objective principle.

Surely, one wants to say, there are cases in which subjective

reasons have a place. I like vanilla ice cream and so does my friend, Fred. Suppose there is only one brick left in the local grocery store at the time we arrive. I have a reason to buy it and so does Fred and to recognize this is to grant Nagel's point about the universality of reasons - I have a reason because I like vanilla ice cream and Fred is in the same boat. It is difficult to see why we should even look for any state of affairs here which is objectively valuable and which both of us, as well as anyone else, has reason to promote. The state of affairs in which I end up buying the ice cream seems no better or worse a candidate for the one which is objectively better, so if there is something to choose between them on such impersonal grounds, it must be because one or the other falls, perhaps quite contingently, under some other description. One could, I suppose, opt for some objective principle such as that the person who has gone without vanilla ice cream for the longest time or the one whose health would benefit most from eating it is the one whose buying the ice cream is objectively valuable, but this seems like an extreme solution... Indeed it looks very much like an attempt to turn the problem into a moral one. This is no accident, as I will argue at the end of this section.

What is it that in Nagel's view makes the use of subjective principles "unacceptable?" It is that:

... a subjective practical principle does not permit one to make the same judgments about others that one makes about oneself, or the same judgments about oneself viewed impersonally that one makes about oneself viewed personally. The application of subjective reasons involves a dissociation of the two standpoints, and a breach in the conception of oneself as just a person among others (p. 116).

The problem, in a nutshell, is supposed to be that using subjective principles leads to practical judgments which have or lack motivational content depending on whether one views oneself as a "someone," impersonally conceived, or as a "me" in one's full concreteness. A subjectivist's judgment to the effect that some person *p* has a reason to do something has no motivational content since it cannot explain any of the subjectivist's actions or desires. But the same is true of the subjectivist's judgments about himself viewed impersonally as just a person (who he happens to be) since, insofar as he views himself impersonally, he is just like anyone else and has reasons in just the way they do. Yet, as we have seen, first personal practical judgments do have motivational content. To judge that one has a reason to do something is to accept a reason to do it while to judge that someone else (or oneself conceived impersonally) has a reason to do something is, for the subjectivist, merely to note a fact. The practical judgments a subjectivist makes about his own reasons will have motivational content if he views himself personally and they will lack motivational content if he views himself impersonally.

Nagel's conclusion from this is that the subjectivist must either deliberate in the personal mode failing to see himself as just one person among many, or else he must remain perversely detached from his active nature:

[I]t is possible to imagine an individual fully capable of occupying the impersonal standpoint and possessing a conception of himself as just another of the world's inhabitants, who nevertheless remained from this standpoint split off, detached from

his practical concerns and his rationally motivated actions (p. 123).

This dilemma is predicated on a particular view of the role of practical principles in deliberation and it is here that I wish to enter a caveat.

For Nagel, practical deliberation is a matter of deducing practical judgments from practical principles. The objectivist does not have the problem of accounting for the appearance of motivational content in his first person practical judgments since the principles he begins with correspond to the assignment of objective value to states of affairs and all of his practical judgments are essentially judgments about who can do what to promote those states of affairs. His will is engaged, so to speak, from the beginning. This points out a certain peculiarity in the notion of motivational content when it is applied not to first-person but to impersonal or third-person judgments. These latter judgments have motivational content on Nagel's account even though they cannot actually motivate anyone to do anything until further judgments about one's situation are added. The impersonal judgment has motivational content in that it embodies the acceptance of some attribution of objective value to something. This acceptance starts the agent's motor running as it were but the clutch is not engaged until he sees what he can do to promote that which has the objective value. The subjectivist, on the other hand, needs to take account of his position in any situation in order to be able to say what state of affairs he has reason to promote, and it is difficult to

see how the bare shift of viewpoint can introduce motivation. There is an important truth in what Nagel says about the subjectivist's "inability to draw fully-fledged practical conclusions about impersonally viewed situations." It points out perhaps the crucial difference between objectivist and subjectivist practical reasoning, but the question remains whether or not this difference points, as Nagel claims it does, to a "practical solipsism" (an inability to see oneself as one person among others equally real) on the part of the subjectivist.

According to Nagel, the subjectivist begins with subjective principles (i.e., ones which make reference to the agent) and therefore his practical judgments about who has what reasons, as long as they are impersonally stated, get no grip on his motivations. Some personal premise must be added in order for the reasoning to result in a first person practical judgment. The problem Nagel sees for the subjectivist here is that it is not easy to see how the addition of a premise, such as "I am the man in the green hat," can introduce motivational content into the reasoning. I suggest that the problem is, however, not a real one since it arises from Nagel's assumption that the subjectivist actually gets to his first person practical judgments by following the process of deduction from principles. Without this assumption (or at least the weaker one Nagel explicitly makes that the subjectivist must be able to get to those practical judgments in this way) the problem of how the motivational content gets into the conclusion does not arise.

Later I will be examining the amoralist's mode of practical

reasoning more closely, but let me just outline here, by way of a digression, an alternative view of practical principles for the subjectivist. Consider someone for whom health (his own) is important and suppose that, being a subjectivist, he assigns no objective value, in Nagel's sense, to his being healthy. He does not think that anyone else need recognize any reason to promote his being healthy, just because he does. His reasoning on some occasion might be reconstructed as follows: "I value my health. Doing A will keep me healthy. Therefore I have a reason to do A." To what does this bit of reasoning commit him? It does not commit him to the view that his being healthy or anyone else's being healthy has objective value in Nagel's sense, but there is still the universality requirement on reasons to be taken account of. If Nagel is right, and I think he is, the subjectivist must be able to accept a practical principle which is universal in form and applies to anyone. Such a principle would be:

(p,A) (If R applies to A, then p has a reason to promote A) where R is a predicate which identifies a reason. Nagel seems to think that the appropriate candidate for R here would be "promotes p's health" and that would certainly yield a subjective principle, viz. "Anyone has a reason to promote his health." But it is not the only possibility. The subjectivist as I conceive of him (the amoralist) would take R to be "promotes the health of an agent, p, who values his health." This predicate also yields a subjective principle, namely, "Anyone who values his health has a reason to promote his health." Now that prin-

ciple will probably seem almost trivial and in a sense it is, but the reason that it seems trivial is that we tend to think of "acting for a reason" as nearly equivalent to "acting to achieve a valued result." Nevertheless the difference between these two ways of understanding R is very important because construing R in the second way relativizes reasons to a person's values. Indeed the personal egoist I talked about in the last chapter counts, on Nagel's definition, as a subjectivist while I classed him as an objectivist. This discrepancy arises because, for Nagel, a subjectivist may allow that some state of affairs has "real" value for a person independently of what he actually values, whereas for me someone who does this is countenancing objective (albeit relativized) values. This need not cause confusion, however, and it remains true that a subjectivist in my sense is also a subjectivist in Nagel's although the converse does not hold. In any case, Nagel can be answered on his own terms and that is all that is required. I will do that presently but let me first conclude the digression on which I have embarked.

In the case above where R is "promotes p's health" and in which the corresponding principle is "Anyone has reason to promote his health" Nagel construes the subjectivist's reasoning as follows:

1. Anyone has reason to promote his health.
- ∴ 2. I have reason to promote my health.
3. Doing¹ A will promote my health.
- ∴ 4. I have reason to do A.

There does seem to be a problem here about where the "motivational content" enters. However, this does not represent the reasoning process of the subjectivist amoralist. If we reconstruct the reasoning using R in the second way above (where R is "promotes the health of the agent, p, who values his health"), we get:

1. Anyone who values his health has reason to promote his health.
2. I value my health.
- ∴ 3. I have reason to promote my health.
4. Doing A will promote my health.
- ∴ 5. I have reason to do A.

Here the first premise seems superfluous since the argument seems perfectly cogent without it. Indeed it seems superfluous in roughly the way a principle of inference such as *Modus Ponens* is superfluous to a deductive argument. It does not figure in the argument but legitimizes it from without, as it were. There is no mystery about where the motivational content comes from - it enters with premise 2 which can hardly be a mere noting of a fact about oneself.

Premise 1, which we have been calling a practical principle, functions as well as an explanatory precept to which the subjectivist commits himself by conducting his own deliberations in the way he does. It does not function as a part of his reasoning, but rather as the principle by which his reasoning is to be understood and hence as the way in which he must understand the reasoning of others if he is to see them as engaged in the same sort of project as he is. (Later I will offer a more complete account of the connection between reasons and values.) This

principle and the subjectivist's acceptance of it, then, constitutes his acceptance of the full and equal reality of other persons. This concludes my rather long digression into the relationship between Nagel's and my version of subjectivism and the way in which the subjectivist amoralist can avoid Nagel's charge of practical solipsism. I now return to answer Nagel directly on his own ground.

Nagel is right in his claim that the subjectivist cannot deliberate to fully-fledged practical conclusions *while maintaining a purely impersonal view of himself*, but that is not sufficient to show that he lacks a conception of himself as just one person among many. In a passage quoted above, Nagel admits that it *is* possible that someone should possess a conception of himself as "just another of the world's inhabitants" and that he should, "*from this standpoint*," remain detached from his practical nature. But I see no need to insist that one always take the impersonal view of oneself nor even that one be able to take it during deliberation. There are things which one *cannot* do from the impersonal point of view, at least if I understand what it is to take up that point of view, even if one operates with objective principles. One cannot perform a deliberate action, for example, as long as one is conceiving of oneself as merely one person among many, because to conceive of oneself in this way is to ignore or bracket those facts about oneself and one's situation which make action possible. One needs to conceive of oneself as engaging the world in one's complete uniqueness. To *have* a certain conception of oneself

it is not necessary for one to be able to use that conception all of the time.

Nagel recognizes, in fact, that "action necessarily involves the personal standpoint. One cannot in general decide what to do unless one knows something about who one is" (p. 121). Thus it is not *action* but *deliberation* which Nagel supposes must be conductable from the impersonal standpoint, if dissociation is to be avoided. But why this difference between action and deliberation? Why should it seem more damaging that a person cannot act without knowing who he is, than that he cannot deliberate without knowing what he cares about? Nagel says:

The fact is that neither of the two standpoints [the personal and the impersonal] can be eliminated from our view of the world, and when one of them cannot accept the judgments of the other, we are faced with a situation in which the individual is not operating as a unit. Two sides of the idea of himself, and hence two sides of himself, are coming apart (p. 119).

Notice the language here. What is it for one standpoint not to be able to accept the judgment of another standpoint? This suggests that the problem is supposed to be that there are judgments which can be made from the personal standpoint which do not find their counterparts in the impersonal standpoint. But this is not the case. All of the subjectivist's practical judgments made from the personal standpoint do have their impersonal counterparts. It is just that the impersonal ones lack, in Nagel's terms, motivational content. The two ideas of oneself are not "coming apart": rather they are just not operating simultaneously. A similar failure of simultaneous operation is

apparent in actions as well, as we have seen. In acting one does not simultaneously see oneself as "just someone" and as "me as I am," but there is no hint here that two ideas of oneself are "coming apart."

What Nagel needs to show, and what he cannot show, is that there is some inconsistency in the judgments made from the impersonal vs. the personal standpoint. It will not do to note simply that one of these sorts of judgments has, while the other lacks, motivational content. The difference between my judgment that "Greg Durward has reason to A" and my judgment that "I have reason to A" is similar to the difference between my judgment that "Greg Durward loves p" and my judgment that "I love p." There is no suggestion in the second case that I am coming apart or that my self-conception is defective. It is just that in the former of the pair, I am reporting on, but not feeling, something I am *expressing* in the latter.

For the objectivist things are different. From the impersonal standpoint his concerns do remain engaged. But that does not show that his self-conception is any "healthier." To carry on the above analogy it is as though his "expressions of love" were reports on something about p ("p is lovable") so that whatever relation he has to p is not merely personal but impersonal. In Nagel's words:

If one acknowledges the presence of an objective reason for something, one has acknowledged a reason for *anyone* to promote or desire its occurrence - at least to desire it, even if he is not in a position to do anything about the matter. This is because objective reasons represent the values of occurrences, acts, and states of affairs themselves, not their values *for anyone* (pp. 119-120).

For the subjectivist all values are values for *someone*.

My contention then, as regards Nagel's main argument, is that he has failed to show that there is anything very peculiar or untoward in the subjectivist's concept of himself which results from the fact that he is unable to deliberate to fully-fledged, motivationally-charged practical conclusions from the impersonal point of view. He cannot, but that is simply a consequence of the fact that he does not assign objective values to states of affairs. The subjectivist can deliberate only in the personal mode, but he can take an impersonal view of his world and his place in it, while he is not deliberating about what to do. In no sense does this show a tendency for two self-conceptions to come apart in a problematic way. All it shows is that, because he does not suppose that anything is valuable quite apart from anyone's actually valuing it, in viewing the world in abstraction from his own actual concerns the world will appear entirely void motivationally. And that seems to me to be quite a sane way for things to be.

It will be recalled that my project in this chapter has been to investigate some arguments which have been advanced in favour of the moral agent's outlook or against the amoral agent's outlook. I have spent a good deal of time discussing Thomas Nagel's thesis for two reasons. First, it seems to me to be a very interesting and sophisticated attempt to show that the sort of objectivist outlook which I identified in the first chapter as one of the central features of moral thinking, is in some sense, required. Secondly, the discussion of his argument has

enabled me, I hope, to explore to some advantage one fairly precise statement of the distinction between objectivism and subjectivism. In the next chapter, I will be looking further into this distinction, but before proceeding there are some loose ends from the foregoing discussion which should be tied off.

I said earlier that the objectivity thesis (that the only acceptable reasons were objective ones) has the effect of making all practical problems seem like moral ones. The reason should now be clearer. The use of objective principles corresponds roughly to viewing every situation from the moral point of view, insofar as one is constrained to look at every situation in the light of an attempt to find a possible outcome for which there is an objective reason. It is not enough to find an outcome such that for each individual person there is some reason or other for him to promote that outcome; rather it is necessary to find an outcome which is supported by the same reason for everyone. It must be an objectively valuable outcome even to people not directly involved.

Finally, some relevant issues are raised in the following passage from Nagel:

The primary opposition to my view comes from egoism.... Egoism holds that each individual's reasons for acting and possible motivations for acting, must arise from his own interests and desires, however those interests may be defined. The interests of one person can on this view motivate another or provide him with a reason only if they are connected with his interests or are objects of some sentiment of his, like sympathy, pity, or benevolence.

... It should be noticed how peculiar egoism would be in practice; it would have to show itself not only in the lack of a direct concern for others but also in an inability

to regard one's own concerns as being of interest to anyone else, except instrumentally or contingently upon the operation of some sentiment (pp. 84-5).

There are three comments I want to make on this passage. First, it is not clear that Nagel has given an account of how motivations can arise except from interests or desires, *in some sense*. It seems that the assignment of objective value to some state of affairs or the acceptance of an objective principle involves taking *some* interest or coming to have *something like* a desire. But perhaps the point is that the assignment of objective value does not itself rest on some interest or desire (Not even to be rational or moral? Even Kant recognized a love of duty and rationality).

Secondly, the thesis of egoism offered in this passage is a general view about the nature of human motivation and should not be confused with any form of egoism I treated in the last chapter. Nagel's argument against egoism (in his sense) would, however, tell against the ethical egoist, the personal egoist and the subjectivist egoist because his argument purports to show that even the sort of objective (relative) value judgments the former two can allow are not enough to avoid dissociation. The archetypical (non-egoistic) amoralist who does not assign objective value to anything is not committed, any more than the others, to Nagel's general egoistic thesis in this passage since he can perfectly well allow that some people do, and others do not, take a direct motivating interest in the welfare of persons generally. There is no reason why even the archetypical amoralist cannot take a direct interest in the welfare of *certain*

persons, although his concern does not generalize to any appreciable extent beyond those persons.

Nagel is involved in showing that the general egoistic thesis he refers to is false and that altruism in his sense is possible and I have some sympathy with his project. It does seem to me that many (even most) people do succeed in assigning objective value to states of affairs and do manage to deliberate pretty much in the way Nagel suggests, using objective reasons and so on. What I am concerned to show in this chapter is that there is no conceptual barrier to someone's thinking in a very different way. But Nagel's attack on egoism does raise the question of whether or not a subjectivist amoralist can adequately understand people who do reason the way Nagel thinks is the only acceptable way. I would be worried if it turned out that he could not. I have suggested that what Nagel calls subjective practical principles are better understood as explanatory precepts which the subjectivist amoralist uses to understand his own and other people's practical thinking. Can these principles provide the subjectivist with a way of understanding objectivists?

The subjectivist cannot admit that any state of affairs *has* objective value, but he *can* understand what it would be like to think this way. He may think there is something wrong with thinking there are objectively valuable states of affairs (depending on the details of his understanding of the concepts involved) or he may just regard it as simply a different way of thinking, albeit one which he does not use. In either case

he could accept the principle (the explanatory precept):

(p,A) (If R applies to A, then p has a reason to promote A),
 where R is the predicate "is a state of affairs to which p
 assigns objective value." As we have already seen this does
 not commit him either to assigning objective values himself
 or promoting any states of affairs whatsoever. His judgments
 about what reasons other people have are ethically neutral,
 which is not to say that he must be attitudinally neutral about
 what other people do. He may on some occasions very much
 want someone to do what they have reason to do and at other times
 want the reverse. Obviously, this makes his use of reason
 judgments open to misunderstanding since we normally use state-
 ments like "You have (he has) a reason to do A" to counsel or
 indicate our desire that someone do something.

Let me summarize very briefly the results of this chapter.
 I suggested that there are two main lines of approach in attempts
 to show that there is something wrong with someone's taking up
 the amoralist outlook. First, it has been suggested that the
 amoralist must be emotionally defective, either because he is
 not endowed with the usual sentiments of sympathy, benevolence,
 etc., or because he lacks the faculty of conscience (or other
 forms of moral insight). Secondly, there are arguments which
 purport to show that the amoralist is irrational or conceptually
 deficient in some way.

Against "sentimentist" theories I have suggested that it is
 not at all clear how emotional responsiveness can underwrite

the whole of moral thinking. Even if sympathy, for instance, can account for a concern for others, how can it make sense of the concepts of right, duty, obligation, etc.? Furthermore, it is very difficult to show that anyone who does not feel very much for the welfare of others in general is actually defective and not merely different. Intuitionist theories are beset with familiar and very serious difficulties.

"Rationalist" theories, on the other hand, seem unable to show either that the objectivist posture evident in moral agency is required of rational agents or that it is irrational or conceptually odd for someone to be little concerned with others' well being. It is possible to make sense of practical deliberation on a subjectivist account of the nature of values, and none of the formal requirements of rationality seem to involve the necessity of deliberating from the moral point of view. Nagel's attempt to show that an integrated self-conception requires the adoption of objective principles fails. The subjectivist and the objectivist differ, to be sure, but Nagel does not show that there is any conceptual stress on the former.

IV

OBJECTIVISM AND SUBJECTIVISM

1 Objectivism

In the first chapter, I offered a characterization of the moral agent in which I concentrated on what seem to me to be his two central features: his recognition of practical reasons which are grounded in the welfare of other persons, and his acceptance of the objectivity of moral requirements. This suggests that there are two dimensions along which to explore the notion of amoral agency. In Chapter II, I discussed egoists since they are clear cases of persons who are amoral in virtue of their disregard for others. It remains to investigate the objectivity/subjectivity dimension, and this will be the task of the next three chapters. Hereafter, then the terms "amoral agent" and "amoralist" should be understood to refer to persons who are not moral agents in virtue of their adoption of an extreme subjectivist posture.

The terms "objectivism" and "subjectivism" have been used by moral philosophers in many, often conflicting, and often not very clear ways and more than once the call to abandon them has been sounded. But, as is often the case, the very fact that the same terms keep reappearing in debates which concern some of the most difficult and fundamental problems in ethics suggests that the problem, whatever it is, will not go away simply by our

refusing to talk about it in the terms which so many people have found it natural to use. As a first approximation, the issue over the objectivity of morality is whether or not, and in what sense, there are things which are "really" right and wrong, morally speaking. Can moral judgments be true or false? Are there moral truths, truths about how things ought to be? When two people disagree over some moral matter, must at least one of them be mistaken? Are the obligations and duties which most people recognize "really" binding?

Two recent authors, unfashionably perhaps, take these questions very seriously. Both J.L. Mackie (1977) and Gilbert Harman (1977) suppose, as I do, that there is a real problem bound up in the notion of moral objectivism, and that the denial of objectivity to morals presents a challenge to morality which cannot be taken lightly. Mackie argues that values (and he includes moral values here) are not objective in the sense of being a "part of the fabric of the world" (p. 15). This does not mean, of course, that there is nothing of value in the world - that the world is grossly deficient - but rather that "there do not exist entities or relations of a certain kind, objective values or requirements, which many people have believed to exist" (p. 17). The issue for Mackie is thus essentially an ontological one, although he does say at one point that another way of expressing the claim that there are no objective values is to say that value judgments do not have a truth value, i.e. they are neither true nor false.

Mackie notes that certain sorts of value judgments, ones in

which there is implicit reference to some fairly determinate set of standards of evaluation, do have a truth value determined by whether the object they evaluate satisfies the standards or not. But then his thesis becomes that those standards themselves lack objective validity. In other words, whether or not value statements posit or assume the existence of objective values directly or indirectly through the operation of conventional standards, makes no difference; there are no objective values. Harman also understands the objectivity issue to involve the truth of moral judgments and he puts the thesis of objectivism in terms of the existence of moral facts.

It is useful to formulate objectivism more broadly than either Mackie or Harman do in order to include as objectivist, theories which deny that moral judgments can be, strictly speaking, true or false. Let us use the term "cognitivist" to refer to theories which hold that moral judgments have a truth value, and the term "non-cognitivist" for theories which deny this. The idea is to formulate objectivism as a theory which is consistent with both cognitivist and non-cognitivist interpretations. Thus, having a truth value will be just one way in which a judgment can be said to be objective. Let us say that an objective judgment is one which can be expressed in a form devoid of indexical expressions such as "you," "here," "now," etc. and which can be assessed, or is in principle subject to assessment, by reference to standards which have validity independently of any particular person's beliefs, attitudes, preferences, and so on. It is important to note that it is the

standards' validity which is to be independent of these subjective factors, and not the judgment itself. The judgment that Sally is angry, for example, is an objective one just in case there is some standard for assessing it which is valid independently of anyone's "subjective states." I assume that the standard of assessing this sort of judgment as for other ordinary empirical judgments is, roughly, correspondence with the facts, the real world, the way things really are, or something like that. It is the acceptance of the existence of an "external world" which underwrites the concept of the objectivity of empirical judgments. We suppose that the standard of empirical truth (correspondence with the facts) is independently valid in the way required.

Naturalistic metamoral theories are ones according to which moral judgments are objective and can be translated, without loss of objective content, into judgments which refer to and assume the existence of only what we normally call empirical properties and relations. A theory which holds that moral judgments are reports of the speaker's emotions, for example, is a naturalistic theory even though such a theory is often classed as subjectivist on the grounds that it refers to the speaker's subjective states. By "subjectivism" I will mean just the denial of objectivism. On a naturalistic theory the standards for the assessment of moral judgments are identically the same as those for ordinary empirical judgments and the methods of science are the appropriate ones by which we are to conduct the assessment (at least if we understand "science" broadly enough). Naturalist

theories, then, are in my terms objectivist ones.

Non-naturalism is the view that while moral judgments are objective, they are not translatable into non-moral judgments because moral judgments refer to or assume the existence of, properties and relations of a special, ethical, non-empirical kind. The clearest example, if not the clearest theory, is to be found in the non-naturalism of G.E. Moore (1903). Plato's theory of forms is another, but more difficult, example. In non-naturalist theories, moral judgments have a truth value and the standard of assessment is correspondence with the moral facts.

I said that non-cognitivist as well as cognitivist theories could be objectivist and I would like to briefly outline a theory of the former sort. Taking our lead from the work of Franz Bretano (1969) and Everett Hall (1961), consider a theory in which moral judgments are taken to express emotions or attitudes, but do not report on them. Moral judgments assert nothing but they can be assessed, on this account, according to whether or not the emotion or attitude is appropriate or correct. The idea is that there are valid standards of appropriateness for the sort of emotions and attitudes expressed by moral judgments. What attitude is appropriate in a given situation will depend on the ordinary empirical facts of the case since it is those facts which distinguish cases in the first place. There is no need on this theory to postulate any special ethical properties which depend (somehow) on the empirical facts, and this gives it one advantage over non-naturalist theories. Because of the crucial role of empirical facts, the theory can begin to explain

why moral language has the apparently fact-stating form it does, and the fact that moral debate centres as it does on getting the facts agreed on by all parties and attempting to assess their relevance. The problem with the theory, however, lies in trying to understand the concept of immutable standards of appropriateness which dictate what responses are the right ones to various situations. Aside from the epistemological difficulty of how we know what the standards are (Brentano thought the emotions come with a kind of felt correctness), there is the question of the metaphysical status and source of the standard. Similar difficulties are to be found in Samuel Clarke's notion of the fittingness of actions to the circumstances of their performance, (Clarke, 1965), but my primary purpose in developing the outlines of the theory is just to show that there can be theories which are non-cognitivist but objectivist.

There is one other form of objectivism which deserves mention as well, since it traffics in neither truth nor the appropriateness of emotions. Beginning with Kant, a kind of theory we might call non-theological imperativism has developed according to which there are objectively valid imperatives which, while not God-given, set standards for, or constraints on, action. Usually these imperatives are grounded in some conception of rationality, so that anyone who is rational is bound to compliance. Actions are assessed according to their consistency with the imperatives. It is possible to distinguish two types of theories within the objectivist imperativist tradition - a cognitivist approach and a non-cognitivist one. A cognitivist version might most naturally

take the form of the claim that ordinary moral judgments are to be understood as stating that some action conforms to or violates an objectively valid moral imperative. A non-cognitivist version might hold that moral judgments are themselves imperatival in nature and are subject to assessment according to whether or not they express valid imperatives. Kant, to my knowledge, never squarely confronted the question of whether or not moral judgments have a truth value, but he certainly did think that actions could be "really" wrong in a very strong sense.

Kant's Categorical Imperative was supposed to rule out certain maxims of action and hence certain actions themselves. (I should say that I assume that any theory which provides an objective standard for actions provides all that is necessary for the evaluation of moral judgments about actions.) C.I. Lewis' Law of Objectivity (1955, p. 39) was advanced as an imperative and Lewis thought it had important practical consequences. Whatever the details of imperativist theories of this sort, the difficulties in them are the same; either there is a problem in showing that the fundamental imperatives have enough content to warrant calling them moral imperatives at all, or there is a problem in showing that they have the requisite objectivity. I have argued in Chapter III that it is unlikely that it can be shown that rationality can ground any appropriate imperatives. Insofar as rationality can be thought of as generating imperatives at all they are purely formal and are quite consistent with amorality. (For a recent attempt to provide an objectivist imperativist analysis of morality, see Bernard Gert's *The Moral*

Rules, in which it is argued that moral rules are validated by the fact that, given certain (morally loaded) constraints, all rational men would take a certain attitude towards certain sorts of actions (Gert, 1966).)

I do not pretend to have given a crystal clear characterization of objectivism, but I am not sure that it is possible to be entirely clear on this matter. It seems to me that some fairly strong notion of objectivity, such as the one I have tried to expound is essentially involved in the ordinary moral consciousness of most people and this makes an objectivist theory of some sort the only possibility for an accurate descriptive analysis of moral discourse. However, I am not at all sure that any objectivist theory can be found which is at the same time satisfactory on general philosophical grounds.

Something should perhaps be said at this point about another concept of objectivity which often finds its way into discussions of this sort, but which is to be distinguished from the concept I am concerned with. We sometimes describe persons, and derivatively their judgments, as being objective when we wish to point out the absence of prejudice and personal bias in their consideration of some matter. The requirement of objectivity in this sense (objectivity in judging) is one normally placed on moral deliberation, but it is best understood, I think, as a *consequence* of viewing moral judgments as objective in the sense I have been discussing. Bias and prejudice have no place in an attempt to discover and apply impersonally valid standards. A certain amount of objectivity in judging can be embraced as a desirable

thing by an amoralist since it makes perfectly good sense for him to attempt to set aside mere prejudice. The amoral agent may well be concerned to ground his valuations in an accurate understanding of the facts. He may, that is, be interested in ensuring that his valuations do not unduly colour or distort his perceptions, and in this sense he may be concerned to be objective.

2 Subjectivism

Subjectivism is the view that objectivism is false. A subjectivist theory, then, is one which denies there are standards for assessing moral judgments which are independent of subjective facts about the persons making them. Because the defining tenet of subjectivism is a negative one, there are many variant positive claims about the nature of morality which, when conjoined to the negative thesis, yield various forms of subjectivism. Gilbert Harman (1977), preferring the term "moral nihilism" to "subjectivism," distinguishes between moderate and extreme moral nihilism. Extreme nihilists hold that "morality is simply an illusion: nothing is ever right or wrong, just or unjust, good or bad" (p. 11), and so we ought to abandon morality. Moderate nihilists, while agreeing that there are no moral facts, hold that "the purpose of moral judgments is not to describe the world but to express our moral feelings or to serve as imperatives we address to ourselves and to others" (p. 12). Harman thinks, as I do, that nihilism in either form runs contrary to common moral thinking and he ends up defending a form of relativism in

order to preserve the possibility of truth in moral judgments. Most subjectivist theories are forms of moderate nihilism although the term is seldom used by their defenders for obvious reasons.

Subjectivist metamoral theories are, of course, all non-cognitivist since one cannot allow that moral judgments have a truth value without at the same time recognizing some standards of truth, and these standards, I take it, to be standards of truth must be objectively valid. The denial of a truth value to moral judgments is not, as we have seen, a sufficient condition for a theory's being subjectivist.

The classic statement of non-cognitivism is found in the works of David Hume. He argued that moral judgments express feelings and therefore cannot be true or false.

Take any action allow'd to be vicious:
Wilful murder, for instance. Examine it
in all lights, and see if you can find that
matter of fact, or real existence, which
you call *vice*. In whichever way you take
it, you find only certain passions, motives,
volitions and thoughts. There is no other
matter of fact in the case. The vice
entirely escapes you, as long as you con-
sider the object. You can never find it,
till you turn your reflection into your own
breast, and find a sentiment of disappro-
bation, which arises in you, towards this
action. Here is a matter of fact; but 'tis
the object of feeling, not of reason. It
lies in yourself, not in the object. So
that when you pronounce any action or
character to be vicious, you mean nothing,
but that from the constitution of your
nature you have a feeling or sentiment of
blame from the contemplation of it.

(Hume, 1888, *Treatise*, pp. 468-9)

Hume's views have been revived by many philosophers who were discontent with attempts to explain the sense in which moral

judgments can be true or false. Unable to subsume moral judgments under an essentially empiricist theory of truth and meaning, many were led to embrace emotivism - the view that moral judgments are but expressions of action.

Emotivists often write in opposition to the objectivist theory of G.E. Moore and consequently often fail to appreciate the range of alternative theories against which they must argue. Objectivism can take many forms but its main virtue is that it takes very seriously what seems to me to be good evidence that moral judgments in ordinary language are taken to express something true or false (or at least something possessing or lacking a strong form of interpersonal validity). It is the failure to appreciate this evidence that has proved the most problematic aspect of emotivism. Emotivists have, somehow, to explain the objective, cognitive flavour of ordinary moral discourse.

One recent attempt by an emotivist to account for the apparent objectivity of moral language deserves mention here. Roger Scruton (1971, "Attitudes, Beliefs and Reasons") argues that moral judgments express attitudes; they do not express propositions which have a truth value. Moral attitudes satisfy three conditions:

- | | |
|---------------------------|--|
| (1) <u>Universality</u> | "Moral attitudes always aim beyond the present instance to some property or state of affairs as such" (p. 47). |
| (2) <u>Overridingness</u> | "Moral attitudes have a particular sort of authority over a man ..." (p. 47). |
| (3) <u>Normativity</u> | "[An emotivist] must ... take seriously the fact that a man's moral attitudes refer beyond him- |

self; they include a desire for a conformity of attitude" (p. 49). "More generally, moral attitudes are concerned with proposing laws to which everyone ... must conform, in specifiable circumstances.... [A] moral attitude will include a desire to influence people and persuade them to conformity to one's own moral views" (pp. 49-50).

It is largely because of the normativity condition that it is possible (indeed necessary) to have reasons for having moral attitudes.

[I]f I have a moral attitude then I am committed to proposing a law of conduct; as a result I am not at liberty to construe any 'Why?' - question [asked about a moral judgment] as a query about my merely personal desires; nor am I at liberty to think that it does not matter if such a question could not be answered in a way acceptable to others (pp. 75-76).

Furthermore, says Scruton:

However sharp the initial disagreement, no moral argument can proceed without the assumption of some *underlying* agreement on the basis of which one of the parties could be shown to be in error.... [I]t can never be proved that there is an ultimate disagreement about what would count as a good reason for an attitude. In adducing reasons, therefore, the appearance of agreement can be indefinitely maintained. As a result, we might say that the subjectivity of moral judgments ... is 'suspended' in rational discussion (p. 80).

Thus, the idea is that because of the nature of the moral attitude and because discussion can proceed only if we suppose some common basis of attitude, we effectively ignore the fact that attitudes are being expressed and speak as though we are stating facts. Since we argue assuming a basic agreement, we can talk as though the intersubjective validity of our judgments can be assumed.

This approach is very Humean, resting as it does on the supposition that people do agree in attitude, at least at some deep level.

Scruton's defense of emotivism is the best one I have encountered, although a number of questions can still be raised. For example, is there any content to the idea of a consideration's being "a good reason for an attitude" besides its being efficacious in producing or maintaining that attitude in a given person or group of persons? In what sense could someone's attitude be "in error" except that it could be based on false beliefs? There may well be acceptable answers to such questions and Scruton seems to have at least recognized all of the features of morality noted here so far. Of especial interest is his concept of suspended subjectivity which he uses to replace the objectivity of cognitivist approaches.

Scruton's central thesis is that in moral debate we only *seem* to find evidence that people think moral judgments have or can have objective validity. What is really going on in moral debate he claims, is that people suspend recognition of the subjectivity of those judgments. This claim, however, will not withstand close scrutiny. In the first place, it must be allowed, I think, that if we suspend our subjectivism we must in some sense recognize it, unless the possibility of moral debate rests on an epidemic of self-deception. Yet, the average man (not to mention many distinguished philosophers) finds the idea that his moral beliefs are really expressions of his own attitudes quite unacceptable. Secondly, what Scruton describes as the

assumption of underlying agreement of attitude and which he postulates as a precondition of moral discussion seems rather to be a consequence of the belief most people have, that since there are right answers to moral problems, reasonable people with enough information ought to be able to agree on the correct resolution.

What does the subjectivist amoralist look like in Scruton's terms? Although the amoral agent does not subscribe to the notion of moral truth neither does an analysis like Scruton's apply to him. (It may be that some emotivist theory which is inadequate as an account of moral language can be modified to describe the way the amoralist thinks but I leave this question aside until later.) The amoral agent does not have moral attitudes as Scruton describes them. His attitudes, or a subset of them may satisfy the conditions of "universality" and "overridingness" since he may have attitudes which are directed toward properties and states of affairs of certain general kinds and which have an overriding motivational efficacy for him, but they do not completely fulfill the condition of normativity. This is because the amoral agent does not assume an underlying agreement on the basis of which he approaches others to attempt to convert them to his way of thinking. He need not suppose that either he, or a person with whose attitudes his own conflict, is, and can be shown to be, irrational or "in error." None of this prevents his having and acting on a desire to have others share his attitudes, but he need not have and need not act on such a desire. Even if and when he does, he does so without

suspending the subjectivity of his attitudes. The amoral agent is not constrained by one of the important consequences of objectivity, or even suspended subjectivity, in morality; viz. that someone who makes a moral judgment is constrained to argue with dissenters indefinitely until he becomes convinced that his interlocutor is irrational, mistaken, incompetent, cognitively impaired, etc. The amoral agent makes no practical judgments which commit him to this.

3 The error theory

John Mackie supports what he calls the error theory. He says:

If second order ethics were confined, then, to linguistic and conceptual analysis, it ought to conclude that moral values at least are objective: that they are so is part of what our ordinary moral statements mean: the traditional moral concepts of the ordinary man as well as of the main line of western philosophers are concepts of objective value. But it is precisely for this reason that linguistic and conceptual analysis is not enough. The claim of objectivity, however ingrained in our language and thought, is not self-validating. It can and should be questioned. But the denial of objective values will have to be put forward not as the result of an analytic approach, but as an 'error theory', a theory that although most people in making moral judgments implicitly claim, among other things, to be pointing to something objectively prescriptive, these claims are all false. It is this that makes the name 'moral scepticism' appropriate.

(Mackie, 1977, *Ethics*, p. 35)

The view that erroneous, confused and even unintelligible assumptions may be built right into ordinary language is not a new one. It is widely supposed that when the ordinary man attributes

colour to an object he means to say something about it which is strictly speaking false. The view that naive realism is built into our thinking and language and that the view is unacceptable on philosophical grounds is also widely held, I think. Language is theory laden and the theories it bears are not only scientific but also metaphysical and ontological - perhaps also ethical. Sometimes changes in theories have little or no effect on ordinary language, but sometimes large networks of words and concepts come into use, undergo changes or drop out of use. In the majority of cases a change in theory has little effect on ordinary language and we can follow Berkeley's advice to speak with the vulgar as long as we understand our words in the correct way.

The amoral agent can be expected to hold his position partly through his acceptance of some form of the error theory, at least if he holds his position reflectively. It will be useful in trying to understand the amoralist's position to consider briefly two other areas besides morals in which an error theory has been influential, and to try and draw some parallels between them and the error theory here under consideration. George Berkeley made a valiant effort to show, not just that his theory of immaterialism is consistent with common sense, but also that it really is the theory of the ordinary man. It was that relatively small group of people who had become befuddled through their acceptance of a thoroughly incomprehensible notion of substance, and not the common man, whose language embodied metaphysical error. Indeed, he says from the mouth of Philonous:

I do not pretend to be a setter-up of new notions. My endeavours tend only to unite, and place in a clearer light, that truth which was before shared between the vulgar and the philosophers: - the former being of opinion, that *those things they immediately perceive are the real things*; and the latter, that *the things immediately perceived are ideas, which exist only in the mind*. Which two notions put together, do, in effect, constitute the substance of what I advance.

(Berkeley, 1965, *Three Dialogues*, pp. 224-5)

Berkeley's attempt to make of common sense an ally in his denial that physical objects have any real and continued existence outside of their being perceived (whether by men or by God) clearly fails. I would say that the view (materialism) against which Berkeley argues is even more firmly entrenched in ordinary thinking than is objectivism in morals.

It is more plausible to regard Berkeley's metaphysics as suggesting an error theory according to which ordinary language about the empirical, physical world implies materialism, but since materialism involves an untenable metaphysics and is in the last analysis quite unintelligible, we ought to abandon, not ordinary language, but the theory of materialism. That is, on the error theory, we ought to begin to think about things in a new way and hence we ought to begin to mean something different when, for example, we speak of things existing somewhere where we are not, or when we refer to physical laws, or when we ask about causes. None of these need upset any of the actual projects we normally undertake, although it will change their complexion somewhat.

In supposing that common sense and ordinary language are

on his side, Berkeley is rather like a metamoral subjectivist who claims not to see that he stands *prima facie* contradicted by common sense and ordinary language, and who denies any need to explain the objective flavour of moral judgments because he claims it isn't there to begin with.

We may set out some parallels among possible positions in the two debates as follows:

- | | |
|---|---|
| 1. Ordinary language and common sense imply moral objectivism and that theory is acceptable. | 1.* Ordinary language and common sense imply that physical objects have a mind-independent existence and that theory is acceptable. |
| 2. Ordinary language and common sense imply moral subjectivism and that theory is acceptable. | 2.* Ordinary language and common sense imply immaterialism and that theory is acceptable. |
| 3. Ordinary language and common sense imply moral objectivism but that theory is unacceptable. We ought to revise our thinking while retaining morality. Amoralism is not the answer. | 3.* Ordinary language and common sense imply materialism but that theory is unacceptable. We ought to revise our thinking while retaining our belief in the existence of a universe which is independent of our minds. Solipsism is not the answer. |
| 4. Ordinary language and common sense imply moral objectivism but that theory is unacceptable. We ought to abandon morality. | 4.* Ordinary language and common sense imply materialism but that theory is unacceptable. We should embrace solipsism. |

I have suggested that Berkeley actually held a position like 2* but that that position is not very plausible. I offered him something like 3* as being a position which is truer to the facts about "the ordinary man." I know of no one who has argued the metamoral position 2 corresponding to Berkeley's actual stand in the materialism debate. In both cases the first and third

positions are the most common. Position 4 is that of the moral nihilist of the extreme variety. With a number of cautions and qualifications which we will encounter later, it is the position of the amoralist.

Before commenting further I would like to introduce one other sort of error theory from another field. Suppose it is argued that there is an error built into theological discourse. Consider in particular the suggestion that the notion of God which is implied in ordinary theological language is one which is fundamentally confused or even incoherent. Corresponding roughly to the positions above, we have the following:

- 1'. Theological language implies the existence of a God with properties p_i and there is such a being.
- 2'. Theological language implies there is no God with properties p_i (or perhaps better "there is a God with properties q_i " where, to keep the parallel, the properties q_i are unobjectionable in themselves but at the same time not very God-like).
- 3'. Theological language implies the existence of a God with properties p_i but nothing does or could (two positions) exist which has those properties (i.e. the theory of God implied in theological language is unacceptable). We ought to revise our thinking about God while retaining religion.
- 4'. Theological language implies the existence of a God with properties p_i but nothing does or could exist which has those properties. We ought to abandon religion.

Here in the religious case as in the others, we find a debate over first, the implications of an established way of speaking and thinking, secondly, the acceptability of the theory contained in those implications, and thirdly, the appropriate response to the unacceptability of the theory among those who reject the theory. I think it is clear that nothing in general can be said

in favour of any of these four sorts of positions (unless to remark that it is always wise to avoid the heroic course of attempting to defend what is really a revisionist thesis by misinterpreting the facts of ordinary language). The decision will depend on, among other things, the details of the position and the appeal of the alternatives and these factors in turn will vary from context to context.

The most interesting feature of the theological example for our present concerns is that alternative 4' is a fairly close parallel to the case of the amoralist, since it involves the giving up of a whole realm of discourse. It seems to me quite possible that someone might be brought to embrace alternative 4' through a conscientious consideration of the various attempts to provide an analysis of theological language and in particular to make sense of talk about God. Someone could, I think, in a mood of thoroughly respectable intellectual despair over the possibility of finding a way of understanding theology, decide that in all probability no way exists which is consistent with not only the ordinary man's conception of God, but also with any way of preserving the distinctive nature of religion and religious concerns. That is, one might conclude that to carry on using religious language in the same old way while thinking in radically different ways, is utterly destructive of the nature of religion. It is important to note that the antireligious conversion in this case is not to be understood as someone's coming to believe the proposition that God does not exist, where before he believed the contrary. Rather, in the context

of trying to come to grips with the concept of God, the conceptual framework in which talk about God made, or seemed to make, perfectly good sense, has crumbled around him, as it were, leaving not some new doctrine, but a void. It hardly makes sense to suggest to a person in this situation to carry on speaking with the vulgar - the point of speaking in religious terms has been lost.

Now some people are able to tolerate a greater degree of imprecision in their concepts than others. Reflections and worries which throw into question a whole realm of discourse for one person might not trouble another at all. It is just a fact that some people, for example, can function perfectly well (or so it seems) while being prepared to admit that there may be no way of understanding God or religious discourse in general. Indeed it is sometimes put forward that an understanding of the sort atheists seek is not only not to be had, but is not to be sought since "God surpasses all understanding." Certainly, these are muddy waters and I am not qualified to wade deeper, but I do think that no matter what one thinks about the possibility of, and desirability of, seeking after understanding in the religious realm, the person driven to atheism through a pursuit of understanding is intelligible and not to be lightly dismissed as having missed the point of religion.

My real motive for discussing the religious case at this length is no doubt apparent. I think that considerations very close to those we have encountered in the religious error theory bear on the issue of amorality. If ordinary language does

imply moral objectivism and if there are grave difficulties, if not demonstrable impossibilities, in the way of an adequate understanding of moral objectivism, then the crucial questions bearing on the reasonableness of abandoning morality become whether or not a subjectivist theory can be found which does not undermine morality in the process of revising our thinking about it, and whether it is reasonable to hold out (in something like a gesture of faith) on the assumption that, whatever the apparent difficulties, an acceptable objectivist theory will be found, if only we keep trying. These questions are extremely difficult and I shall not pretend to offer definitive answers to them. My central project, after all, is to attempt to give a characterization of the amoral agent and to show that it is not altogether an unreasonable response to certain reflections on the nature of morality to decide to abandon moral thinking altogether. I am not quite as ambitious as Berkeley.

4 Against objectivism

It is probably impossible to show conclusively that no theory could be found which both adequately analyzes ordinary moral language and thinking, and at the same time is deemed completely satisfactory on general philosophical grounds. Perhaps no philosophical theory is ever completely satisfactory, but there are certain considerations which incline me to think that a modicum of added skepticism may be warranted here. However, the most I can do at this point is to survey briefly the most prominent candidates for the source or grounds of the alleged objectivism

in moral judgments and to suggest in each case the problems which must be overcome. Anyone who has a detailed theory at hand will likely find the discussion far too skeletal to be truly worrisome but I can hope at least to indicate the area in which further debate must centre.

(i) Empirical truth as the source of objectivity

A naturalist theory is one according to which moral judgments can be translated, for purposes of assessing their acceptability, into ordinary empirical judgments of one sort or another. Sometimes one finds attempts to effect the translation by the device of equating ethical properties (goodness, rightness, etc.) with some non-ethical properties (conducive to hedonic experiences, tending towards the maximal satisfaction of human desires, etc.). Sometimes the translation is between whole utterances (e.g., "You ought to ..." becomes "In our society there is a rule requiring that people ..."). In all naturalist theories moral judgments are translated into statements which can, in principle, be assessed by the ordinary standards of empirical truth and the method is roughly the method of "science."

One general problem faced by all naturalist theories is that none of the translations seem to capture the meaning of the original moral judgment. Of course there is no need to absolutely insist that they do, but then there is the question of just what the relationship of the original moral judgment to the translation comes to, if it is not an identity of meaning. If the notion of a translation is not the appropriate one, then it may be possible to argue that the objectivity of moral judgments is grounded in

the truth of corresponding empirical statements while remaining uncommitted to the claim that moral judgments are strictly speaking true or false. The idea would be that a moral judgment is acceptable or correct if (and, presumably, only if) some empirical statement is true. Now a dilemma presents itself. Either the nature of the empirical statements involved explains the relationship we have noted between moral judgments and motivation, or it does not. If it does not, one fundamental aspect of morality remains deeply and disturbingly mysterious. If it does, it seems this must be because those empirical statements report facts which are intrinsically motivating.

Suppose then that a certain moral judgment is correct or appropriate if and only if some empirical statement is true. Anyone who accepts the truth of the empirical statement is thus rationally bound to accept the moral judgment. But to accept a moral judgment sincerely is to accept a reason and a sufficient motive for doing something. What sort of fact can be counted on to operate in this way? I have already discussed what seems to me to be the most sophisticated attempt to establish the connection between facts and the motivational content of moral judgments in my treatment of the ideal observer theory.

The basic problem is that any attempt to ground the objectivity of moral judgments in the empirical facts tends to ignore the fact that people are not actually motivated by the same things. This means that we must give up the interpersonal validity of moral judgments or disqualify certain people's reactions to the facts on some grounds or other. The difficulty

in the latter is to avoid loading the dice illegitimately.

(ii) Non-empirical truth as the source of objectivity

Even less satisfactory are theories which posit the existence of odd ontological entities, be they objective values or non-natural goodness. J.L. Mackie (1977), directs virtually all of his anti-objectivism arguments against such theories.

I have nothing to add to the chorus of voices which has been raised against non-naturalism. If Occam's razor ought to be applied for the excision of unnecessary entities, a much less delicate instrument is required to eliminate unintelligible ones. I should say, however, that if I could bring myself to accept and perhaps understand such entities, I expect I would have a very good analysis of morality indeed. The appeal of a Moorean analysis which takes ordinary language very seriously can be largely accounted for by two things. First, moral language is indicative in form and obeys the ordinary logic of indicative sentences, and this makes the categories of truth and falsehood natural ones to apply. Secondly, common sense suggests that the way to understand the notion of truth is in terms of correspondence between judgments or sentences or propositions and "the way things are" (the facts). What could be more in harmony with the ordinary man's use of moral language than a theory which holds out the possibility of just the sort of correspondence he assumes is possible in other realms of discourse?

The most promising line of development for a concept of non-empirical truth probably lies in the attempt to generate a general theory of truth which has some form of correspondence

theory as a special case for empirical claims, but I cannot adequately deal with that idea here.

(iii) Universality of sentiment as the source of objectivity

I have already had something to say in Chapter III about theories which posit the universal existence of a moral sense in persons and I have nothing to add here. The idea that we come equipped, as it were, with some special ability to discern through some emotional response the correct standards of action for human beings is not very plausible. The evidence for the existence of anything like a moral sense is very sketchy. When the influence of other possible sources of what may appear to be the operation of a natural and authoritative moral sense are taken into account, the residue seems unimpressive.

Much more sophisticated forms of the sentimentalist theory can be formulated. For example, it might be argued that it is in the very nature of man to live socially and it is not only to be expected that people will learn to care for the woe and weal of their fellows in ways which are expressible in the requirements of morality, but that the only truly human existence possible to man consists in his leading a life in which these sentiments exert their influence to a given extent. Thus, it is not just that most people happen to share certain emotional or attitudinal responses to certain situations which is important for providing the standards for action. Rather it is the fact that these emotional responses represent the core of being human or at least an important part of the core. To be subject to the requirements of morality is to be subject to the promptings

of one's truly human nature. The standard for assessment of moral judgments is their conformity with the prompting of an ideal human emotional set. The theory needs to be fleshed out considerably, of course, but the crucial elements are the positing of a truly human sort of emotional makeup and the claim that it is that makeup which explains both why moral judgments have the motivational efficacy they do, and the idea that everyone is at bottom subject to the same objective requirements.

There is certainly something appealing about this sort of approach but it relies on ignoring or downplaying the valuational character of the notion of being "truly human." The attempt to ground morality in some conception of an ideal human nature is interesting, but it must be kept in mind that what one gets is not strictly speaking an objective morality unless it can be shown that the ideal is an ideal independently of anyone's embracing it. One might be forgiven for suspecting that there is as much in our truly human nature which is destructive to morality as there is to support it.

However persuasive descriptions of an ideal are, I do not see how it can be shown that it has any validity other than the relative validity it may have for persons who actually accept it as their own ideal. Even the notion of health, broadened perhaps to include such things as emotional health, spiritual health and psychological health, rests on either the concept of function or of purpose, and it is difficult to see how it can be argued that the emotional or attitudinal responses and sensitivities of the healthy person are those which underlie

morality, especially if one concentrates on the requirement that the conceptions of function or purpose must be valuationally neutral or at least universal.

I do not doubt that there is some conception of an ideal human nature at the bottom of most people's thinking about morality but I do doubt that it is possible to make very good sense of the idea of that ideal's objective validity.

(iv) Universality of subjective valuations as the ground of objectivity

If it could be shown that everyone does, or has reason to, value something or some set of things and that that something either is intrinsically moral (e.g., the maximization of human desire satisfaction) or could be attained only through the adoption of moral values (i.e., values which are such that anyone pursuing them is automatically a moral agent) then it could perhaps be maintained that moral judgments are objective in the sense that *anyone* who understands what he has reason to do would adopt a moral outlook or perhaps even certain moral principles. If, for example, morality necessarily pointed the way to a state of affairs in which everyone would be better off than under any other arrangement, and if anyone always has overridingly good reason to do what will make him better off, then it would be plausible to argue that it is in the very nature of human existence that moral thinking finds its source.

This approach could take at least two broad forms. Either it could be held that there "really" are reasons which are independent of what people actually count as reasons (i.e., some

things just are reasons and some are not, independently of what people find important) or it could be held that as a matter of fact (perhaps of very deep fact) people all "really" do agree on what counts as a reason or would if they had enough information. In Chapter II, I discussed the attempt to ground morality in a universal egoistic concern and found it wanting; yet self-interest seems to be the best candidate for a universally shared propensity. Arguments that people have "real" concerns which may be quite inconsistent with what they actually seem to concern themselves with, face the problem of giving content to the idea of an unvalued value.

There is very good evidence that people do have pretty much the same sorts of concerns and that there is a relatively small class of values in terms of which virtually all human action can be understood. There is nothing particularly surprising in this, but neither is there anything very helpful to those who would attempt to ground morality in it. Even if people value pretty much the same sorts of things, it would be necessary to show that certain values *ought* to take precedence over others, that values *ought* to be aggregated in certain ways, and so on. Any concerns which are universal are very general and the relationship among them is far from universally uniform. But even if there were a general congruity of valuations, it is difficult to see how such distinctively moral notions as those of duty and obligation can be given a foothold in this fact without being contorted in the process. Obligations and duties have more *authority* in the ordinary moral consciousness than can be

grounded in facts of what people happen to find important.

I have been arguing that it is difficult to find any adequate grounding for a very strong formulation of objectivism. I have not, and could not, prove that the attempt must necessarily fail nor that it has been or must be without value. And even if strong objectivism cannot be given unshakeable foundations there may well be theories which can support the objectivist posture of moral agency. Whether these theories turn out to be strictly speaking subjectivist ones, in my terms, is not, of course, the important issue. What matters is the impact they have on moral concepts and ways of thinking. In Chapter VI, I consider the possibility of someone's accepting a (strictly) subjectivist analysis of moral concepts which is consistent with the objectivist flavour of moral discourse.

5 The problems in metamoral subjectivism

I have suggested that objectivist metamoral theories, however well they reflect the form and meaning of ordinary moral discourse, are subject to criticism on the grounds of a general philosophical concern for clarity and intelligibility. I should like now to turn to an examination of some of the problems which must be faced by subjectivist theories. In this section I concentrate on theories which purport to give a non-revisionist account of morality, and in the next section I will deal with theories which see some problem with ordinary moral thinking and language and which are essentially revisionist in spirit.

If I am right in claiming that the objectivist form of moral language is not merely a surface phenomenon but rather reflects a deep acceptance of some kind of objectivism on the part of moral agents, it is to be expected that the main problems with metamoral subjectivism will centre on its failure to provide an adequate interpretation of that fact. A subjectivist must provide not only an explanation of the objective flavour of moral language, but also an interpretation of it in terms which are consistent with the denial of the objectivist thesis itself. I have discussed Roger Scruton's attempt to do just this using the concept of suspended subjectivity and I have argued that his account fails. Since I think that his theory is the best attempt to defend emotivist subjectivism against the charge of inconsistency with ordinary language, I will add nothing further here.

There is, however, another form of subjectivism which has received widespread attention and which deserves comment. I discussed earlier the strain of objectivist imperativism which derives from Kant, but there are also subjectivist versions of the imperativist approach. R.M. Hare (1963) argues, for example, that moral judgments are essentially prescriptive but that there is no objectively valid standard which exists for assessing them which is independent of individual persons' desires, preferences, beliefs, etc. Hare speaks of moral agency in terms of "playing the moral game" and he recognizes the possibility that not everyone will or even need play the game. (Playing the game, as Hare conceives of it involves using moral language according

to certain rules.) Ordinarily, I think, it is supposed that while someone may not play the game of morality, not playing involves making a mistake in some sense. Another way of putting the point is to say that we ordinarily suppose that everyone is subject to the demands of morality, that everyone ought to play the game whether they do or not - in some sense it is not optional whether one is in or out. But further, according to Hare, even in the case of those who enter the moral arena, there are no objective standards for evaluating the moral pronouncements of participants. Moral judgments can be assessed only by reference to reasoning rules which operate on facts about persons' "subjective states." Hare says:

The rules of moral reasoning are, basically, two, corresponding to the two features of moral judgments which I argued for in the first half of this book, prescriptivity, and universalizability. When we are trying, in a concrete case, to decide what we ought to do, what we are looking for (as I have already said) is an action to which we can commit ourselves (prescriptivity) but which we are at the same time prepared to accept as exemplifying a principle of action to be prescribed for others in like circumstances (universalizability). If, when we consider some proposed action, we find that, when universalized, it yields prescriptions which we cannot accept, we reject this action as a solution to our moral problem - if we cannot universalize the prescription it cannot become an 'ought'.
(Hare, 1963, *Freedom and Reason*, pp. 89-90)

The essential objection to Hare's analysis is that it allows in principle unresolvable moral disagreement. Of course it is no argument against a metamoral theory that it fails to show how actual disputants in a moral debate can come to an agreement, but it is an objection to a theory that it leaves no room for the

notion of one person's contradicting another's moral judgment in cases where both parties are engaged in serious moral discussion following the rules of the moral language game. That is, an adequate descriptive metamoral theory must allow the possibility of there being, and ought to provide some way of understanding, genuinely contradictory moral judgments. Hare's analysis does allow that moral judgments can be *conflicting* in roughly the way that commands can be conflicting but I think that ordinary moral thinking requires more than this.

According to Hare, moral judgments can be criticized along two dimensions. First, if a person who makes a judgment does not embrace the universalized judgment that applies to *anyone*, then he is required to withdraw the judgment or to admit that it is not a moral judgment at all. Secondly, if the person does not embrace the prescriptive implications of his judgment in the full range of actual and hypothetical cases, then again he must withdraw his original judgment, as a moral one. No doubt this does provide a good deal of leverage on persons who are concerned to defend themselves morally, but there is some question whether it goes far enough. Hare himself revealingly chooses to call persons who pass the tests proposed but who do so only by being prepared (hypothetically) to endure the painful personal consequences of their moral principles (putting their "ideals" above their comfort as it were), "*fanatics*." That suggests a recognition that there is actually more to playing the moral game than Hare sometimes allows (in particular that it involves displaying a concern for the welfare of others).

What would be required to turn the bare bones thesis of "universal prescriptivism" into a viable metamoral theory?

First, I think that a concern for the welfare of persons must be somehow built into the moral language game. Secondly, the methods of criticizing moral judgments must be in principle tight enough to ensure that in "morally charged" situations there is a fairly narrow range of acceptable actions. Finally, moral requirements must be binding on everyone independently of whether they choose to play the moral game or not.

Hare actually goes some way toward developing a theory along these lines in the chapter on Utilitarianism in *Freedom and Reason*. In dealing with the question of what is to count as being prepared to embrace the universalized prescriptions entailed by one's moral judgments he says:

For if my action is going to affect the interests of a number of people, and I ask myself what course of action I can prescribe universally for people in just this situation, then what I shall have to do, in order to answer this question, is to put myself imaginatively in the place of the other parties (or, if they are many, of a representative sample of them) and ask the same sort of questions as we made the creditor ask when he had imagined himself in the situation of his debtor. And the considerations that weigh with me in this inquiry can only be, How much (as I imagine myself in the place of each man in turn) do I want to have this, or to avoid that? But when I have been the round of all the affected parties, and come back, in my own person, to make an impartial moral judgment giving equal weight to the interests of all parties, what can I possibly do except advocate that course which will, taken all in all, least frustrate the desires which I have imagined myself having? But this (it is plausible to go on) is to maximize satisfactions (p. 123).

If to test one's moral judgment involves imaginatively taking on everyone's position (ignoring one's interests, inclinations, etc. except to count them as someone's), then anyone should, if the exercise is a conceivable one at all, arrive at the same answer. The maneuver, then, effectively satisfies the first two requirements the bare bones theory failed to fulfill.

There is still the matter of entering the moral language game, however. That we ordinarily suppose our moral judgments to apply to anyone regardless of whether they are prepared to enter moral debate (provided of course they are moral (not non-moral) agents) is fairly clear I think. Hare, however, has very little to say on this point. Because what he offers is a method of assessing the moral judgments of persons who are willing to make them, in terms of their sincerity, sensitiveness, and informedness, the results of the assessments will get a grip only on those persons. The judgment Jones makes that Smith ought to pay his debts then is relevant to Smith only if Smith is prepared to enter the debate according to the rules of the game. If Smith offers the judgment that he ought not to pay Jones we have a moral debate, and Hare's fleshed-out theory holds some promise of there being a determinate resolution. The resolution will have the form of the discovery of some principle which Jones and Smith and any other moral agent will adopt. The resolution will not, however, involve the discovery of a principle which has objective validity since the only kind of validity it has, it has in virtue of the (subjective) acceptance by the parties involved of the rules of moral reasoning. What is required to

attain objectivity in moral judgments (and I think the lack of objectivity signals a failure in the theory to describe moral thinking the way we find it operating in the real world) is some demonstration of the necessity of participation in the moral debate. It is this last step which seems insurmountable.

6 Revisionist theories

No doubt some readers will be puzzled at the stress I have placed on the issue of objectivity, either because they regard it as obvious that my objectivity requirement is too strong and that I have gone too far in attributing a central status to it in my characterization of ordinary moral discourse, or because it seems a simple matter to correct the error by which strong objectivity has come to be a part of our prereflective moral consciousness. It is my experience that most people can be brought to recognize the kind of objectivity I have been discussing in their own moral thinking and that they react rather strongly to the suggestion that things might go on pretty much as before even if they give up that feature.

Let us assume in this section that there is a significant objective element in moral language and thinking generally, and consider the possibility of construing the work of subjectivists as revisionist in spirit. Although most subjectivists have not subscribed to what I have called the error theory, a few, including R. Robinson (1948), Mackie (1977), and Hume, on some interpretations, have. Revisionist theories are much more difficult to assess than descriptivist theories since what will

constitute an acceptable revision to ordinary thinking will depend on such things as how great a change is required, how the purpose of moral thinking is conceived, and so on, and there may well be individual differences along these dimensions. The most I can hope to do here is to suggest some of the features and functions of morality which seem to me to be underwritten by the notion that moral judgments are objective and therefore which can be expected to founder to a greater or lesser extent should a subjectivist analysis be adopted. Then I will examine briefly some ways which have been suggested to avoid these subversive tendencies within a subjectivist framework.

The assumption of objectivity in morals has a dual importance. On the one hand, it holds out the possibility of strongly justifying individual actions and, more importantly, social customs, laws and institutions. The idea that individuals have "real" objectively grounded obligations and duties to one another and to a "morally legitimate" government, is a powerful social tool. The ability to fall back on some notion of objectivity in these matters lends an important force to attempts to keep society "on the rails," so to speak, when some of its members violate the expectations of the populace at large. A "You ought ..." (as an authoritative but impersonal demand) is at the same time easier to impose and easier to accept gracefully than a "We want" The ability to appeal to the objective values of liberty, peace, and justice is an importance social force.

On the other hand, the idea that there are "really" right and wrong modes of behaviour combined with the difficulty of

knowing with certainty just what *is* right and wrong creates an atmosphere, ideally anyway, in which rational, disinterested, and as far as possible dispassionate, discussion appears as the appropriate method for resolving differences in the moral realm.

The assumption of objectivity also affects the atmosphere of debates in the public forum. If political and quasi-political questions are at least in large measure supposed to be dealt with by considering the moral issues involved, the influence of power and interest groups can be expected to be minimized or controlled. Under ideal circumstances the resolution can be advertized as the best attempt of the best minds to determine the requirements of morality and everyone can be expected to accept the results as the conscientiously determined best approximation to the correct solution. Contrast this with the view that there is no correct solution to political questions but only solutions which various groups and individuals prefer.

Finally, the assumption of objectivity tends to make people feel that they are accountable to others for their actions (at least those which affect others). If the standards by which actions are to be assessed are objective, then virtually anyone is in a position to call for a justification of anyone else's actions. There is a presumption that in matters of fact (whether moral or empirical) one ought to be able to convince any reasonable person that one is correct or at least not clearly wrong. In any community there will be widely held views on the right- and wrong-making properties of actions, and it will be generally very difficult for people to find arguments which they can expect

others to accept for views which run contrary to public opinion.

Thus the idea that there must be a publicly defensible justification for actions in the moral realm creates an atmosphere of public accountability which will seem natural and non-arbitrary to everyone. Other persons will ideally appear as allies in the deliberative process. Advice can be sought and advisors, supposing there to be a correct solution in principle to the advisee's problems, will have a reason to set aside their personal stakes, if any.

Revisionist theories, then, must not only offer a philosophically satisfactory non-objectivist value theory, but they must also show that the acceptance of subjectivism is consistent with moral language and thinking pretty much the way we know it. The possibility that morality cannot survive the abandonment of objectivity has seemed real enough to many philosophers. Arguing that emotivism discredits morality, H.J. Paton says that, even if people's approvals would not necessarily change were some form of emotivism to be adopted:

... the so-called "moral" actions would have no value other than that of satisfying my contingent desires, and so ... they would not be moral actions at all (p. 121).

I can see little hope for the spread of "moral" action in the world unless we are prepared to accept and to act upon a law which we believe holds for all men alike and bids us treat others, not simply as a means to the satisfaction of our own contingent desires, but as ends in their own right. To accept and to act upon such a law as binding upon all men is to adopt in action a principle which is objective, not merely in the sense that it treats myself and others impartially, but in the sense that it springs from a reason which

necessarily manifests its impartiality in acting (as in thinking) and is distinct from any merely contingent desire.

(Paton, 1948, "The Emotive Theory of Ethics," Symposium. *Logical Positivism and Ethics*, pp. 122-3)

This passage is interesting for what it does not quite say. Paton, of course, is an objectivist and he is arguing that subjectivism is not only false but pernicious. The task of morality would be undermined if people came to be emotivists and actually started to view their moral beliefs as expressions of non-necessary (merely contingent) desires or emotions. If morality is to thrive (if people are generally to perform the sort of actions we generally suppose to be required morally) then people must be objectivists. The conclusion has to be that anyone who is concerned that moral actions be done must attempt to see that people in general retain the outlook of objectivism whatever its intellectual credentials.

Interestingly there are also internal pressures which induce a natural drift to objectivism. Moral agency and objectivism are mutually reinforcing. Suppose that someone is an objectivist about values in general. Then, in the absence of peculiar circumstances, such a person can be expected to suppose that other persons similarly concerned to find out what is really of value will be in agreement with him. The search, by a group of persons, for a value scheme to which all can ascribe is likely to issue in precisely the sort of values we call moral (as well, perhaps, as some others). Because values are practical, people will tend to insist that their own welfare is valuable, not only to them, but in itself and will be led to adopt a

theory according to which the welfare of persons *per se* is what is valuable.

Conversely, and more to the present point, people who are moral agents in that they have typically moral concerns, deliberate from the moral point of view, and so on, will tend to use objective-sounding language and will come to think that their way of thinking is not one they have simply contingently adopted. A moral agent could hardly exist in a world in which others did not also think and act similarly, and the sheer fact of congruency among the values and behaviour patterns of such a group of persons would soon extinguish any linguistic conventions which made practical judgments seem clearly personal. Being prepared to deliberate from the impersonal, moral point of view, already takes one far enough that the remaining step to the view that one is discovering what ought to be done rather than merely deciding what to do will seem quite natural. Moral thinking puts the question "What do I want?" into the background in favour of the question "What shall I do?" Because moral deliberation involves addressing this question from a point of view which deemphasizes any purely personal desires and preferences, the moral agent will see himself as simply a "someone." From this standpoint the objectivist problem "What ought to be done?" is barely distinguishable from the question "What shall I do?"

Finally, the moral agent sees himself as just one person among many and is already concerned with the interests of others sufficiently that he will find an objective value vocabulary quite comfortable. Because he is prepared to subtract from his

view of himself, in his practical thinking, much of what individuates him, he will easily find himself supposing that what is important to him is actually important objectively.

In conclusion, then, I suggest that whatever the ultimate philosophical verdict on the intelligibility and acceptability of objectivism, the bulk of mankind is bound to accept, however unreflectively, the thesis of objectivism. Not only is this a natural course in the ways outlined but, as Paton argues, morality cannot survive a general, reflective and sustained subjectivism.

In Chapter VI, I will be discussing the possibility of individual (versus general) moral subjectivism but I hope enough has been said to show that the prospect of treating subjectivism as a revisionist theory directed at the general public is one not likely to succeed and that prospect of finding a satisfactory, objectivist metamoral theory is slim, but not entirely non-existent. In the next chapter I will attempt to explain, in more depth than I have been able to manage so far, how a reflective subjectivist amoralist might think (practically) and how his practical vocabulary might differ from that of a moral agent.

V

REASONS AND VALUES - A SUBJECTIVIST OUTLOOK

1 Facts, reasons, and the definition of "moral"

If subjectivist amoralism is to constitute a real, interesting, and rational human possibility, there must be a view of the nature of practical reasons which is tenable and consistent with this outlook. Otherwise the amoralist could not hold his position reflectively and rationally. Obviously, the problems involved in developing an acceptable theory of practical reason are very difficult ones and I cannot hope here to do the subject justice. However, I will attempt to produce the outlines of such a theory in an effort to show that there are no obvious and compelling objections to taking the subjectivist amoralist seriously.

Peter Singer (1973) has recently suggested that moral philosophers have expended far too much energy on attempts to define "the moral." The important question, the issue that really matters, according to Singer, is how statements of fact are connected with reasons for acting. To show the ultimate futility and practical insignificance of the debate over the nature of morality, he distinguishes what he takes to be the two extreme metamoral positions.

Neutrality is the view that a principle is a moral principle

for a person if it has an overriding practical force for him. A moral principle, on this view, may have any form and any content whatsoever. Thus, the personal egoist's principle, "I ought to pursue my own self-interest" is a moral one, even though it is not universal and takes no account of others' welfare. *Descriptivism*, at the other extreme, is the view that moral principles must satisfy certain conditions of form and content. A theory which identifies moral principles as ones which are universalizable and concerned with the welfare of persons, is an example of a descriptivist theory.

The advantage of the neutralist position according to Singer is that it leaves no troublesome gap between someone's acceptance of a moral principle and his acting in accordance with that principle:

The neutralist is able to explain why, if a man acts on the basis of a coherent set of principles at all, he will act in accordance with his moral principles. If a man recognizes that a certain action is prescribed by his overriding principles, he surely will do that action, if he can (p. 52).

There is, however, a problem with the neutralist view. It is that many of the sorts of facts which we normally take to be relevant to moral problems can be ruled out as irrelevant by someone whose overriding principles make no room for them. The egoist, for example, can ignore with impunity, on the neutralist view, facts about the damaging effects of his actions on others' welfare.

At the other extreme, descriptivism has the advantage of ensuring the relevance of certain familiar forms of arguments

and of certain sorts of facts to moral discussion. The problem with descriptivism is that it fails to ensure a connection between a man's moral principles and his actions:

We are not, on the descriptivist view, free to form our own opinion about what is and what is not a moral principle; but we are free to refuse to concern ourselves about moral principles.... So morality may become irrelevant to the practical problem of what to do (p. 53).

Descriptivism allows us to argue from facts to moral judgments but leaves us free to ignore those moral judgments, while neutralism ensures that moral judgments are relevant to action but leaves us free to adopt anything at all as a moral principle.

Singer considers a "middle position" between neutralism and descriptivism according to which moral judgments are necessarily prescriptive and universalizable. But neither the extreme positions nor the middle position can succeed in bridging the gap between facts and action. He concludes that:

... there are limits to what any account of morality can do. No definition of morality can bridge the gap.... It follows that the disputes over the definition of morality and over the "is-ought" problem are disputes over words which raise no really significant issues (p. 56).

It is worth noting that Singer does not even consider the possibility of combining the elements of neutralism and descriptivism in the way I have. I have suggested that a moral judgment is one which both has a certain form and content and is overriding. Of course, some will see it as a disadvantage to my account that it leaves open the possibility of someone's not making nor being committed to making moral judgments, but that is

different from the problem Singer sees in descriptivism. Descriptivism allows that someone might sincerely make a moral judgment and yet fail to concern himself with it when it comes to acting. My account insists that this is not possible (except in peculiar cases such as those in which we invoke the notion of weakness of will). Of course, my account does not bridge the gap between facts and action either if that is taken to involve showing that the facts commit a person to a moral assessment and that the moral assessment commits him to acting in a certain way. But my account does begin to close this gap for those who are prepared to make moral judgments. Someone, for example, who accepts the objectivity thesis and who is committed to deliberating from the moral point of view may be constrained to interpret the facts as providing a justification for a very narrow range of possible actions.

By putting the debate over the nature of morality in terms of a dispute between neutralists and descriptivists, Singer is able to make that debate seem trivial. Of course, a neutralist and a descriptivist *might* disagree in a merely verbal way as, for example, when both of them hold a crude form of emotivist theory according to which moral judgments are expressions of emotional responses to various situations. The neutralist emotivist might allow that any sufficiently strong and overriding emotional response can generate a moral judgment while the descriptivist might prefer to count only certain very general emotional responses directed at human welfare as moral. In such cases one might expect the neutralist and the descriptivist to use

pretty much the same sort of "arguments" in any real-life practical situation, although the terms in which considerations are advanced might differ somewhat. Their emotivism will ensure that their approach to practical problems will be very similar.

But there is no guarantee that the differences between neutralism and descriptivism will always be so slight. Contrast the neutralist emotivist above with a descriptivist objectivist who believes that there are true and objectively valid moral principles which are universal in form and which concern the welfare of all persons equally. I think it is clear that in a real-life situation these two persons can be expected to argue in very different ways. Just as one example, the neutralist emotivist may resort without any qualms to pure rhetoric and exhortation in order to convince someone to do something, while the descriptivist objectivist may make appeals to the self-evidence of the truth of various basic moral requirements. No purely verbal agreements will bring these two into the same camp.

What Singer largely ignores is the fact that neutralism and descriptivism are positions which are advanced, not out of the blue, nor even as purely descriptive accounts of the nature of moral judgments. Rather they are positions which are held largely as a result of the attempt to get clear on just the issue Singer sees as the important one - the relation between statements of fact and reasons for acting. Neutralists typically suppose that the gap between facts and action must be closed by each individual's own choice of practical principles. Descriptivists

typically are impressed by arguments which purport to show that certain forms of principles and certain sorts of facts have a special claim on the actions of rational human agents.

I can see little reason to think that attempts to understand the nature of morality must generate mere verbal disputes, nor that philosophers have in fact been distracted from important issues by their concern to understand moral discourse. Neutralists don't *just* prefer to use the term "moral" the way they do; rather they prefer to use it that way because of some deeper views about the nature of reasons for acting. The same is true of descriptivists.

Moral discourse provides important clues about the way people think about reasons. The attempt to provide a philosophical analysis of moral discourse involves the attempt to evaluate those ways of thinking as well as the attempt to understand them. The terminological differences between neutralists and descriptivists should be understood as symptomatic of much deeper differences.

As I have noted, my characterization of the moral agent does not lie at or between the extremes Singer identifies. According to my analysis, the amoralist's failure to use moral language signals much more than a verbal difference. The moral agent's language embodies a view of the nature of practical reasoning which the amoralist rejects. I suggest, *contra* Singer, that the attempt to give an analysis of moral language and principles is capable of highlighting rather than obscuring the questions he sees as the important ones.

In the rest of this chapter I shall be attempting to give an account of how a reflective amoralist might think about practical reasons in his own case. That is, I am interested to see, in outline, how an amoralist might understand what's going on when he acts for reasons, deliberates about what to do, and offers explanations and justifications of his actions. I have said something about how a reflective moral agent views these matters and if the moral agent, as I have described him, is right, then the amoral agent is making a mistake somewhere. But I have also suggested that no conclusive arguments in favour of the moral agent's outlook have been offered and this leaves us free to examine alternatives. I am not committing myself, nor need the reflective amoralist commit himself, to the claim that there is ultimately only one acceptable or true account of the matter. It may be that there are simply various views available, none of which can, even in principle, be better argued for than another. I see no reason to feel embarrassed if our conclusion is that there is no clearly correct answer to such questions. The outlook of the moral and the amoral agent are inconsistent, of course, and this puts limitations on their tolerance to the views of the other. But each may recognize in the other a coherent position which can no more be proven correct or false than can his own. This chapter, then, attempts to outline one view an amoralist might take of the nature of practical thinking.

2 Reasons and explanation

The notion of a practical reason has two faces - one turned toward explanation and the other toward justification. Consider first the role that the appeal to reasons plays in explaining a person's actions. In order to explain some person's doing something by pointing to his reasons it is necessary that the reasons adduced should be that person's reasons in a fairly strong sense. While we sometimes say that someone has a reason to do a particular thing whether or not he knows it (i.e. whether or not he is aware of the facts in which the reason resides) no reason of this sort can explain his action. One thing we can require is that for something to count as a reason which explains a certain act, it must be something which played some role in bringing the action about. That does not mean that reasons must function strictly speaking as causes, nor that a person must be aware of what his reasons were in order for his actions to be explainable in terms of his reasons. It does mean, however, that there must be a true description of the production of the action which makes references to what are adduced as the person's reasons for acting as he did.

In order to make any progress with the question of how reason explanations work it is necessary to say something about what it is for someone to do something for a reason. Unfortunately this is one of the most difficult areas of ethics, in my view, and I cannot hope to do more here than to gesture toward a theory. To begin to appreciate the problem one need only reflect on the various sorts of ways of completing the sentence

"His reason for doing that was"

Following Thomas Nagel (see Chapter III, section 5), let us identify reasons with predicates which apply to actions. That is, let us assume that, given an action A which a person P has done for certain reasons, it is possible in principle to locate predicates R_i which P believed applied to A, and which can be adduced as the reasons P did A. As we have seen the amoralist recognizes only reasons which carry a reference to the agent concerned (subjective reasons, in Nagel's terms), but more than this, he grounds all reason attributions in particular facts about the agent (for example, his actual concerns or even his acceptance of objective values).

To do some action A, then, for reasons given by the predicates R_i is, roughly, to count the applicability of R_i to A as considerations in favour of doing A. This idea could be generalized to take account of cases in which someone does A because he has reasons for not doing not-A, but we can safely consider only the simpler case of positive reasons. The main difference between the moral and the amoral agent, again using Nagel's terms, is that the former will be prepared to back his subjective reasons with objective ones (i.e., with ones which make no reference to the agent). The moral agent may realize that the amoralist is not prepared to do the same, but he will not condone this since he thinks that, at least in some cases, the only ultimately acceptable reasons are objective ones.

In general, of course, there are many reasons which together explain a person's action. Here, as when we offer causal

explanations of some event, we tend to ignore all but the most salient features of the situation. Which of a person's reasons we offer as the reason(s) will depend on such things as what we suppose our interlocutor to know already about the circumstances and the purpose for which we are attempting to give an explanation of the action.

How do we go about trying to say what a person's reasons were on a particular occasion of action? As a first approximation consider the claim that: a person's belief that R applied to A was one of his reasons for doing A if he would have been less likely to do A had he not believed that R applied to A. This comes fairly close to what is needed, I think, but there are problems with it. First, there are cases in which it is true that someone would have been less likely to do A had he believed that not-R applied to A but where this does not show that P was counting on R as a consideration in favour of A. The belief that R applied may have been instrumental in producing A by, for example, distracting him in a way which resulted in his being in a new situation in which other quite different reasons were present and sufficient to produce the action A. It is possible to imagine such a case in which the original R would normally function as a reason *against* A. Something more than a mere probabilistic connection between P's believing that R applies to A and the likelihood of P's doing A, is required.

Secondly, the analysis is incomplete in that it provides no way of distinguishing actions done for reasons from other happenings or "mere doings" which are made more probable because

of the holding of a belief about the nature of some behaviour. In other words, we must already know that we are dealing with a case of acting-for-a-reason or we may be led to misidentify as reasons things which operate in some other way to produce action. I am not going to attempt any further development of the notion of "counting something as a consideration in favour of acting" here, but something along the lines of what I have suggested fits well with the amoralist's position.

What I am after, in part, is an account of what it is to act for reasons which is highly general as regards the sort of thing which a person might count as a reason. The task of trying to explain why someone acted as he did is essentially a matter of determining what it was about the action which disposed him to do it. If a situation is one in which the person has consciously raised for himself the question of what to do and has consciously identified the important facts about the situation and about the likely consequences of various acts open to him, then he may be able to provide an explanation of his actions himself, but this is by no means guaranteed or necessary. Someone may act for reasons of which he is not consciously aware.

Correlative with this view of reasoning is an account of values. Insofar as a person's reasoning displays a pattern of systematically treating certain factors as counting or not counting for acting certain ways, one can formulate a description of the person's value scheme - a description of what things are important to him. To some extent such a pattern must exist

for action to be explainable at all since it is only through a recognition of, or presumption of, a pattern which pervades a person's actions that it is possible to say with any confidence what on any particular occasion a person's reasons are.

3 Reasons and justification

As I argued in the first chapter, being a moral agent involves adopting a view of reasons which makes them criticizable externally or objectively. That is, the moral agent supposes that at least some sorts of reasons (the most important sort) can be assessed on some grounds which have a validity for everyone regardless of what values they actually have. Reasons on this view are good or bad, acceptable or unacceptable, valid or invalid, independently of who actually counts what as a reason. This is what makes the moral notion of justification such a powerful one. There are various ways of arguing the underpinnings of this notion of justification.

G.E. Moore (1903) for example, thought that there is a special non-natural property of goodness and that certain sorts of states of affairs have this property and others do not. Those states of affairs which have it are valuable (ought to exist) and those which do not are either neutral or possessed of dis-value. According to Moore, we can somehow gain intuitive insight into just which states of affairs do and which do not have value. The attempt to justify actions and the point of deliberating in the first place both rest on the idea that anyone has reason to promote those states of affairs which contain the most goodness.

Moore's approach is an extreme one but it is instructive in that it shows how powerful is the drive to find a theory which can make sense of the concept of justification which we find in ordinary moral discourse.

The minimal account of reasoning outlined above suggests some "internal" grounds for a critical assessment of a person's reasons. Insofar as an action forms a part of a coherent pattern of action so that it can be viewed as directed toward the achievement of some valued end, it can be appraised as efficient or inefficient. If other behaviour of the same person warrants the attribution of some value(s) to that person then a particular action can be criticized as conforming to, as promoting, or as defeating that value. Actions can be undertaken on less than adequate information given the importance of the goal it is intended to serve, and so on. In each case the evaluation is premised on the values attributable to the person and is internal in that sense. To attempt to show that an action is not deficient in any of these internal dimensions of criticism in which the actual values of the individual are used to ensure the relevance of the criticism to the agent, is to offer an "internal justification" of that action. An amoral agent has no use for any concept of justification which goes beyond this internal critique.

The moral agent on the other hand is committed to the possibility of what we can call "external justification." External justification involves trying to show that an action conforms to some standard of acceptability which has (or ought to have on pain of ... (what?)) weight for anyone, regardless of what happens

to actually matter to him. It is sometimes claimed that the very notion of deliberation brings with it the concept of external justification. In Chapter III, I discussed C.I. Lewis' attempt to make good on just such a claim. The basic idea is that the performance of an action, especially a deliberately undertaken action, involves one in supposing that it is an action which has reason behind it, or at least that there is something to be said for it. Acting deliberately, the argument goes, one must act in a way which one holds to be acceptable, not only for oneself but for anyone. As Sartre puts it, for reasons I do not understand:

Now, I'm not being singled out as an Abraham, and yet at every moment I'm obliged to perform exemplary acts. For every man, everything happens as if all mankind had its eyes fixed on him and were guiding itself by what he does.
(Sartre, 1947, *Existentialism*, p. 24)

There are three directions which seem to be the most plausible from which to attempt to establish the supposed connection between external justification and deliberation. The first, represented by Lewis, holds that unless the point of deliberation is to get the right, true, objectively correct answer to the question "What ought I to do?", then deliberation has no point. I argued briefly in Chapter III that the practical question which deliberation is to solve can be expressed as, "What shall I do?" and that there may well be a point to deliberation which can be explicated entirely in terms of the aims, desires, values, and so on of the deliberating agent. Action may, if it is to have any point at all, require an internal justification, but it does not need an external one.

The second argument can be put thus: to deliberate is to locate and to give due weight to reasons for and against the various actions open to one. To act as the result of deliberation is to adopt some set of reasons as sufficiently supportive of that action. But this is just to accept a justification for doing the thing in question since whatever is a reason for one person is a reason for any person, so the argument goes. Reasons are impersonally valid because they are formulable so as to apply to anyone relevantly like the agent in relevantly similar circumstances. T. Nagel (1970, p. 65) links deliberation and justification in just this way. The answer to this argument is quite straightforward on the view of reasons I have been taking. It is of course true that anyone relevantly like a person who has a reason to ϕ also would have a reason to ϕ in similar circumstances, but this is trivial if part of the relevant similarity is that one count the same facts as favouring ϕ -ing, since that is just what having the reason amounts to. The most the argument can show is a point about the internal justification of action. The point about external justifications is that they hold for anyone in the sense that an action which is externally justified is so in virtue of some feature of it (perhaps its intrinsic nature or perhaps its consequences) which is such that anyone has reason to promote it or at least not to interfere with its occurrence. To admit that an action is externally justified, then, is to make a judgment which has potential practical consequences for the speaker in terms of what to do vis-a-vis the action of someone else.

The third sort of argument holds that to act deliberately is to make of oneself an example of what man ought to be. Every act, on this view embodies a normative principle of universal applicability. We have encountered one statement of the view by Sartre and he states it again like this:

In fact, in creating the man that we want to be, there is not a single one of our acts which does not at the same time create an image of man as we think he ought to be. To choose to be this or that is to affirm at the same time the value of what we choose.... We always choose the good; and nothing can be good for us without being good for all.

(Sartre, 1947, p. 20)

It is true that to choose to be this or that sort of person is to affirm the value of being that sort of person, but it may be that the value one affirms is the value of being that sort of person *to oneself*. When I take as my ideal the "strong silent type" and conduct myself accordingly I need only be affirming the fact that I value my being that sort of person. I do not need to be affirming (indeed I may deny) that everyone ought to be strong and silent nor that it is good or valuable objectively or absolutely that I or anyone else have that sort of character.

I am not denying that one *can* take the approach rejected by the amoralist and think about deliberation and justification in the way most people are inclined to. Of course one can (although I am not convinced there is anything very compelling to be said in favour of thinking this way). What I am committing myself to is the idea that one need not use the concept I have identified as "external justification" and that in giving it

up one is certainly not placing oneself in the realm of the unintelligible.

Let me summarize the account so far given of the nature of deliberation and add some remarks which may be helpful. Deliberation has as its point the performance of internally justified actions, i.e., ones which incorporate all the available information, which are guided by a coherent and consistent value scheme (if the agent has and is concerned to have one), which embody a "decent" epistemological theory, and so on. Insofar as deliberation is undertaken in a situation in which one's action will, or can be expected to, affect one's realization of a life in which one's values are maximally achieved, the criteria for the assessment of the quality of deliberation are obvious enough. Every deliberate act implicitly claims to satisfy the requirements of internal justification, and so is at the same time a statement of the agent's values. I have said that values can be inferred from how an agent acts if we can reconstruct his reasoning from our knowledge of his beliefs and our hypotheses about his values. In one way, then, values emerge from deliberation; yet in another, deliberation makes sense only if a person has values already. I do not think there is any real paradox in this. We deliberate because something is at stake but in the process of weighting the facts we sometimes also come to appreciate them in a different way than we have in the past. Experience teaches us to make more and more subtle distinctions, and this produces a shift in values. We discover that past goals, once achieved, have yielded more pain than we expected or offered

less challenge or failed to sustain our interest as well as we had expected. That is not to say that there are a small number of things which are our "real" values and that all other concerns are instrumental.

Deliberation is a counting of factors in the situation as for or against acting in various ways but that is not all it is. It also involves acquiring beliefs and modifying them, searching out possible courses of action, reconsidering priorities and values, predicting consequences, experimenting with "ways of being," and much more. Nonetheless, the central activity of all of this is the weighting of the facts in favour of or against various actions.

4 Practical debate and vocabulary

It may seem from what I have said so far that the amoral agent is restricted to receiving and offering judgments which are grounded in his or other persons' existing concerns, at least if he is to deal with others sincerely. But that is not quite true. He can also engage in discussions in which it is the avowed purpose of one or more of the parties to induce a change in the practical basis of some of the others. When the amoralist is attempting to get someone else to come to value something he does not now value, or to disvalue something he now values, he will, of course, do so not because he thinks that person is (objectively) mistaken but because, for one reason or another, he finds it important to effect the change. As I have said before, there is no reason to suppose that there is

necessarily something selfish or untoward in this concern. He may, for example, attempt to instill certain values in his children out of concern for their future well-being and a belief that only people who are concerned about certain things live a life of the sort he is concerned to have his children live.

There are many ways of effecting changes in someone's valuations. One can better acquaint someone with "the facts" in the hope that his values will alter once certain errors one suspects he has made have been corrected. One can also point out that there are logical or practical inconsistencies among his values on the supposition that an attempt to remove the inconsistencies will produce the desired alteration in his values. Another way is to attempt to present an ideal to him in such a way as to "capture his imagination," perhaps by getting him to admire or identify with some fictional or real person who holds or exemplifies the value one is trying to advocate. Other methods include various forms of brainwashing and therapy, as well as more subtle kinds of psychological manipulation. Some of these and a multitude of others are "rational" in the sense that they do or can take the form of a reasoned argument, taking as premises facts about a person's pre-existing values and beliefs; others are non-rational, relying on the knowledge of psychological effects of treating people in certain ways. The distinction between rational and non-rational approaches is not, however, very clear and neither is its significance.

Because much of our ordinary practical language is objective in form (i.e., there is no grammatical difference for the most

part which can serve to distinguish judgments of "fact" from judgments of "value") and since moral judgments imply an acceptance of objective reasons, a willingness to argue in certain ways, and so on, there are problems for an amoral agent who attempts to use ordinary practical language in conversing with others. The problem is essentially one of avoiding some of the implications which can normally be drawn from statements about what people have reason to do, what they ought to do, what they are obliged to do, and so on. According to the account I have given of what it is to have a reason, for example, to say that someone has a reason to do something is not necessarily to express any support for that person's doing that thing nor to indicate an acceptance of that sort of reason in one's own case.

Earlier I discussed briefly the moral and the ethical use of "ought," but there is another use which is non-moral and non-ethical and which can be used by the amoralist. That is the use of "ought" in which, for example, an historian might say that Hitler ought to have invaded England early in the Second World War; that is, in which a person adopts uncritically and hypothetically, the values of someone else and then reasons from that other person's point of view. This is roughly the "ought" of advice since it is the hallmark of advice that unless some indication is given to the contrary, it is assumed that the "ought" offered is offered from the viewpoint of the advisee and is connected with the values he actually embraces. Still, advice is seldom offered sincerely without the advisor's being content with, if not his being actively concerned to promote,

the advisee's doing what he has reason to do. Thus, it is often difficult to see that "ought" is being used in the purely descriptive sense I am suggesting is open to the amoral agent.

None of what I have said about the ability of an amoral agent to use "ought" in its "descriptive" sense should be taken to suggest that he lacks the ability to speak evaluatively. At the very least he can say what he values and why (where "saying why" involves indicating those features of the things valued which are the grounds for his valuing them). Furthermore, given a situation in which a group of persons share and are known to share certain values, ordinary evaluative (objective-sounding) language can be used without being misunderstood, as long as all of the parties recognize that their evaluative utterances are grounded on nothing more than their (perhaps quite contingent) valuations.

5 What can be valued?

So far I have not made any room for a critique of values, except to say that a person may be susceptible to considerations of consistency among his ends and of efficiency in his choice of means. Are there things which in themselves are such that they must be valued by any (rational) person, or things which cannot be valued by any (rational) person? If we assume that any action, character trait, state of affairs etc. can be described in value-neutral terms we can eliminate one kind of example. Clearly there are descriptions which employ words which, if someone is speaking ordinary English, imply (not merely contextually) that

a person has certain values. To describe something as "good," "worthwhile," "interesting," etc. and yet to act in ways which reveal that the things so described hold no importance to one is to have spoken misleadingly. One way of putting this is to say that some words have an evaluative component to their meaning. What I want to do is eliminate, as uninteresting in the present context, claims such as that anyone must (rationally, logically) value or take an interest in good food, interesting pastimes, worthwhile books, etc., since insofar as someone accepts these descriptions of something, he necessarily betrays a valuation on his part. We need not concern ourselves with someone who violates the meaning rules knowingly.

It is not obvious, however, just what descriptions do carry "emotive meaning" and it is a matter of debate, and to some extent, of choice, where lines are to be drawn. Is it a misuse of language, for example, to recognize something as being in one's interests while failing to give it any *prima facie* weight in one's deliberations? In a sense the answer matters little since what we need is just linguistic consensus on when one is speaking misleadingly either (a) because one's *description* uses words with a certain evaluative meaning or (b) because most people in one's linguistic community assume that certain things, however described, are valued by everyone. The very fact that there is generally a broad consensus of values among members of any linguistic community virtually guarantees that the distinction will not be very clear. For example, when A notes that doing X will harm B, are we entitled to infer that he takes himself to

be reporting a reason not to do X, and if so is our entitlement grounded in the meaning of "harm" or in our knowledge of the moral reasoning basis prevalent in his community or perhaps in something else?

Let me assume that it is possible to sort out those descriptions which linguistically entail evaluations since I want to concentrate on those which do not. Let me further assume that a "non-evaluative" description can be given of anything which can be valued. This seems to me quite plausible especially if we switch to reason terminology and put the question in terms of what predicates can, or must, or must not, be treated as identifying reasons for or against various sorts of actions in the deliberations of a rational person.

There is one familiar logical constraint on how a rational agent must reason practically. If a person has the existence of some state of affairs as an end or goal, then the fact that doing a certain thing will bring that state of affairs about must, if it be known to the agent, count as a *prima facie* reason for doing that thing. It will, in general, be only a *prima facie* reason since there may be many ways of achieving an end and since the means of achieving one end may be inconsistent with the means of achieving other ends. It is not always clear whether a given action is done as a means or an end, nor that all actions must be one or the other, but when an agent's ends have been discovered, means/end reasoning is rationally called for.

But are there things ("states of affairs" is the most general description I can think of) which must be valued or which cannot

be valued? We have already encountered some arguments which purport to show that consideration of the interests of others must weigh in the deliberation of any rational and fully human person. But what about the agent's own interests? Can someone be oblivious to his own self-interest?

The answer would seem to be a pretty clear "no," at least if we attend to the long-run. Anyone who pays no attention to his self-preservation will perish unless his life be maintained by others - but even then it is hard to see how a fully human individual can exist in this manner. I am only concerned with persons who are engaged in the world via having projects of some sort, and a human vegetable, kept alive on life support systems in a hospital ward, not just temporarily but as his mode of being, hardly qualifies.

It is to be noted that the fact that all rational agents do concern themselves with their own welfare to some extent, as a condition of their existence, does not in itself show that self-interest is a requirement of rationality. Something more is required to establish that it is irrational not to care about whether one lives or dies. It may be that, except for very desperate cases, any clear-thinking person will realize that there are things which are important enough to him that on balance he has stronger reasons for preserving his life than for relinquishing it. I suspect that that is the case and, if so, we can conclude that virtually all rational people will take an interest in their own well-being, if only as a precondition of their being able to strive toward the attainment of goals

they hold even more dear. But that does not show that an interest in one's own health, happiness, or whatever is a condition of being rational. If someone really doesn't care about anything or cares only to die (and there may be rare instances in which this is so) then it is not irrational for him to cease acting or to act to terminate his existence.

6 Valuing and commending

There are two questions which, in the present context, must be distinguished; first, "What things can be valued by a person?" and secondly, "What things can a person call 'good'?" Calling something "good" involves accepting the meaning rules governing "good" and it involves interacting with other people linguistically and this latter involves one further in accepting contextual rules of meaning. By "contextual meaning rules" I am referring to those conventions and understandings which underlie and make communication possible and profitable, but which are in large measure independent of the meaning of the words used. The context of an utterance affects the meaning of the utterance through the dynamics of the situation.

"Good" is a word in ordinary language and there are meaning rules governing its use, however hard it is to state what they are. It functions in communication in certain ways and it is possible to misuse it and to use it misleadingly. Thus, one is not free to call "good" anything one pleases. This, however, in itself, is not enough to show that one cannot find valuable anything at all. Just because it would be misleading or incorrect

to say of certain things that they are good, it does not follow that it is irrational or incorrect in any other way to treat those things as ends. The context of justification and the context of explanation can be pried apart to a large extent.

Calling something "good" is different from valuing it (or treating it as good, or finding it good) to the degree that one's valuations are abnormal in ways which would invite misunderstandings were one to employ, without warning, the ordinary evaluative vocabulary of one's social group. Using words such as "good," "valuable," "interesting," "important," and many others, places one within the context of justification. Any ongoing social unit will have a fairly stable value framework which structures justification and determines what reasons will be generally acceptable. For the most part, people's reasons for acting will not only explain the behaviour to others but justify it to them; but that is just to note that most people do operate within the ordinary value framework.

Because a society's value framework is not completely determinate and because the weighting of values within the framework will vary from individual to individual there is a region in which people can disagree about the justifiability of actions while agreeing about their explanation. Person A and person B may both accept the same reason explanation for some action of person C but they may disagree about whether there were, and about whether C had, reasons which justify the action. If A and B are typical moral agents they will suppose that in at least some cases there is a question about the existence of an external

justification of C's action.

In her article, "Moral Beliefs," Philippa Foot argues that:

... assumption (1) ["that some individual may, without logical error, base his beliefs about matters of value entirely on premises which no one else would recognize as giving any evidence at all." (p. 84)] is very dubious indeed, and that no one should be allowed to speak as if we can understand 'evaluation' 'commendation' or 'pro-attitude', whatever the actions concerned (Foot, 1958-9, p. 94).

It should be noted that Foot speaks of "beliefs about matters of value" and of "evidence" for those beliefs and in a sense this begs the very question at issue. It may well be, indeed I have argued this as well, that moral agents cannot suppose that matters of value in general are subjective in the way she objects to. But as I have already argued also, evaluation and commendation are public activities carried on with a rule-governed vocabulary. Having a pro-attitude is, I think, a different matter unless it be defined into this technical concept that someone has a pro-attitude to something only if he is prepared to engage in pro-selytism on its behalf. If, however, having a pro-attitude involves little more than behaving fairly consistently to bring about or maintain the existence of the object of the attitude and things of that sort, then I do not see how Foot's arguments justify her conclusion.

Basically, Foot argues that many (all?) evaluative words (eg. "dangerous," as well as "good") have an internal relation to their object, so that not just anything can intelligibly be described by these words. For instance, we would not understand someone who said that clasping one's hands is a good action

because we can see no point to doing so and because:

It is surely clear that moral virtues must be connected with human good and harm, and that it is quite impossible to call anything you like good or harm (p. 94).

There are two points here, one about the concept of moral virtue (or "good" in its moral sense) and one about "good" in general. The second point is the one I wish to concentrate on, since I am prepared to grant the first as according with my own views.

The point Foot is trying to make is, I think, that (even leaving aside the question of justification) we cannot accept reason explanations of just any sort. We simply wouldn't understand someone who valued clasping his hands whether or not he commended the activity to others or attempted to offer purportedly objectively valid reasons for performing this action. Let us consider this case. (We can suppose that the agent realizes that others have different values and that to enter an attempt at justification would be futile.) Can we understand (explain) such an action?

The obvious suggestion is that we can explain someone's clasping his hands three times, say, if we simply note that doing so is something which is important to him; he values doing this thing. The immediate and natural reaction to this suggestion is to ask, "But *why* is it important to him?" In one sense we *don't* understand the action because we don't understand his attribution of value to this activity. There are two ways we can gain understanding, and thus render the proffered explanation acceptable.

One way is by relating the peculiar values to ones we are more comfortable with, by filling in a "special background" (Foot's term). If our person finds the sucking noise aesthetically pleasing or if doing hand clasping is part of a ritual with some intelligible significance, we might be happier with the explanation. But what if efforts to find a means/end or inclusion relation of this sort fail? There is another way of rendering the explanation acceptable, I suggest: it is one we are inclined to leave as a very last resort but it is not different in kind.

Suppose we investigate the hand-clasper's behaviour very carefully and discover that he becomes agitated and unhappy if we prevent him from clasping his hands. He tries to teach his children to clasp their hands and is disappointed in them when they fail to do it the way he wants. Perhaps enough evidence might accumulate so that we are prepared to allow that he *does* value performing this action in itself and that he is not doing it from habit or to achieve any further end. It is important to him. Of course, this would be a very surprising result and we still would not understand why he values what he does. His action is one which is explainable by his values but we want, and do not have, an explanation of his values. To this extent the explanation is partial and we cannot be completely comfortable with it until we find some account of why he thinks as he does or perhaps how he came to be the way he is. Even if we do find such a further explanation, however, this will not cause us to retract our explanation in terms of his values.

If it is admitted that it is logically possible for someone to value clasping hands (perhaps it would help the imagination

to make our person a Martian) then it is logically possible that we should come to share his outlook. In this event, I submit, we would not necessarily have any further explanation of why someone might find such a peculiar thing important but then we wouldn't feel the need either. We would then be in the same situation vis-a-vis the hand clasper as we actually are vis-a-vis someone who actually does share our ultimate (non-derivative) values. To be sure our psychological makeup would be different in the two cases, but logically the situations are identical.

Explaining action by reference to reasons is a two-level process, only one of which is usually apparent. In the first place it involves adducing the reasons or values which a person has which operated to produce his action. In the second place it involves sharing the values and reasons which are adduced. In virtually all cases we are actually confronted with, there is no significant disparity in values and the explanation succeeds. It is possible, or so I have been trying to show, that the explanation should fail at the second level while remaining intact at the first, without anyone's being irrational or making any logical error. Someone could not, of course, have any utterly mysterious values and still qualify as a fairly normal human being, and I should say, to avoid possible misunderstanding, that I am not supposing the amoral agent necessarily has such values.

The important point in all of this is that whatever rules govern public discussion of values and reasons, and whatever

meaning rules apply to evaluative words, it remains possible for someone to actually deliberate and attend to reasons in the most unusual fashion. The amoral agent may have to be very careful about what he *says* if he wishes to avoid being misunderstood even if his values are merely abnormal (perhaps in their relative importance for him) and not downright unintelligible. Language cannot dictate what people can find important, but it can and does embody what people do find important.

7 Facts and reasons again

Let us return now to the question with which we began this chapter; "What is the relationship between statements of fact and reasons for acting?", setting aside, as Singer urges us to do, the question of the nature of morality. The answer according to the subjectivist view I have been developing is simply this: "The relationship, in its logical essence, is whatever people make it."

Ultimately reasons connect with a person's values and it is possible for someone to value anything (at least I don't see how we can ~~eliminate~~ much *a priori*). Complications arise when someone acts inconsistently, treating apparently similar situations in dissimilar ways or apparently dissimilar situations in similar ways, when someone's behaviour generally does not permit us to form a coherent view of his values, and when someone seems to have very strange values. But in some such cases we may simply have to conclude that the person just doesn't have reasons or that he has very unusual ones.

In the next chapter I will be reopening the question of whether or not a recognizably moral outlook is consistent with the subjectivist theory of reasons and values outlined above. But before proceeding I would like to take a brief look at an objection to the kind of subjectivism about reasons being developed here.

Many people find it difficult to make sense of the notion that reasons can be subjective in the way suggested. Only if the validity of reasons (or values) is supposed to reside in something outside of the individual, it is thought, can reasons actually function in the way they do. Thus, for example, existentialists are often accused of presenting a theory which inevitably produces a stultifying despair - a practical paralysis. Even existentialists themselves recognize this to some extent: Sartre calls the reaction to the realization of radical freedom "anguish."

Certainly there is something disconcerting about the notion that we choose our values and reasons freely with no possibility of referring our choice to any ultimate and absolute standard. But whether the uneasiness comes from an absurdity in the very idea of such a free choice is another matter. I think a significant part of the suspicion that there is a logical difficulty here comes from thinking of the creation of values or reasons as involving *choice* in the normal sense. I avoided using the concept of choice above deliberately for this very reason.

I think that normally we think of choosing as something we

do for reasons and it is obvious that this immediately sets up an infinite regress. If we choose our reasons, and do so for reasons, then there must be reasons for our choice of reasons, and reasons for our choice of those reasons, and so on. That is why I prefer to talk in terms of treating facts as counting for or against various actions. It is at least less problematic to suggest that we don't necessarily have reasons for weighting the facts the way we do.

"That's just the way I am," can serve to stop the regress when we steer clear of the notion of choice, but it does so in a non-trivial way. Valuing what we do (treating facts, motivationally, in the way we do) is, in fact, the way we define ourselves, not only to others but to ourselves. At some point, at least if we refuse to pass on the responsibility to some objective realm of values or to our past, we must simply stand up and be counted for what we are. Not "Here I stand I can do no other," but "Here I stand I *will* do no other."

Of course, someone may still object that we cannot really take responsibility for who we are, even as regards our fundamental values, because our character and personality have been determined by our past. I can hardly hope to resolve the free will/determinism debate here but there seems to me to be two ways of getting around this objection. First, it should be noted that there are at least two senses of "responsible" which must be distinguished. To say that A is responsible for some state of affairs (where A may be an event, state of affairs or a thing, person, object, etc.) may be to say that A was the

cause (or an important contributory causal factor) in the production of B. Alternatively, where A is a person, it may be to invoke some quasi-legalistic or moral convention operative in the community whereby persons like A in certain circumstances are *held* responsible (*made* liable and so on) for B's occurrence.

It is a part of our ordinary moral view that no one ought (*prima facie*, at least) to be held responsible for anything they didn't at least help to bring about. But that requirement is a moral one and not an analytic one and in any case holds only *prima facie*. We do hold people responsible for things they could not have been expected to prevent, although not very often. Sometimes superiors in an organization take responsibility for the actions of their subordinates *carte blanche* and they are fired because of the misdeeds of those subordinates. This practice does not rest on the view that the superior was (causally) responsible for hiring (or not firing) incompetent subordinates but on something like the usefulness of having the practice in force. If the possibility of *taking* responsibility and of being *held* responsible were confined to cases of *being* (causally) responsible, no such practice would exist. Parents are held responsible for the actions of their children, very often when it would be quite fanciful to suppose that parents have the amount of control over their children's behaviour which would be required to say that they caused the behaviour or even could have prevented it.

The idea of taking responsibility for what one has not brought about, thus, is a familiar enough one that it is not

utter nonsense to attempt to rescue the possibility of one's taking responsibility (and being held responsible) for one's character even under the hypothesis of determinism. This view could be made more sophisticated and perhaps more acceptable by distinguishing among various sorts of causal factors and providing a rationale for treating some of these as generating the possibility of personal responsibility and others not. The second possibility, of course, is the libertarian denial of determinism, perhaps in Sartrean form, but I do not propose to discuss that further here.

The main point I have been trying to make is that there is no absurdity or obvious conceptual difficulty in supposing that people can choose or be responsible for their values in some way which is consistent with the belief that those values have no objective claim to be adopted. Someone can recognize that his values lack objective validity without being plunged into a nihilistic despair. Someone can see himself as the ultimate source of his own values and in that sense accept responsibility for being the person he is. The belief that values are not chosen for objective reasons does not necessarily reduce one to conceiving of oneself as a being whose nature is determined from without nor does it necessarily produce despair over the meaning of life.

This concludes my attempt to provide the outlines of a description of how an amoralist might think and speak. A great deal more would have to be said to transform the above into a fully articulated theory of practical reasoning. I have not

tried to carry out this full development but it seems to me that there is nothing obvious which suggests that it could not be done.

VI

MORAL AND AMORAL SUBJECTIVISM

1 The individual moral subjectivist

In Chapter IV, I argued that an acceptance of the notion of objective values and an essentially moral outlook are mutually reinforcing and that there is a "natural drift" toward moral objectivism. If my conjectures in that regard are correct then we can expect that any attempt at an analysis of ordinary language will seem adequate only to the extent that it makes room for some fairly strong sense in which moral judgments can be objective in nature. There is a problem, or so I have argued, however, in finding an objectivist metamoral theory which is not only adequate to the data of ordinary moral discourse but which is acceptable on general philosophical grounds.

I have not shown that no acceptable objectivist metamoral theory can be found and I have not ruled out the possibility that individual persons could resist the appeal of objectivism while remaining moral agents. It is this latter possibility which I wish to examine in the first part of this chapter. Let me call such a person a "moral subjectivist." I expect that anyone whose moral concepts are correctly analyzable along subjectivist lines will hold his position fairly reflectively, if only because it runs against the more natural objectivist

view. Thus there is little harm, I hope, in referring to him in terms which suggest that the position is a theoretical one.

If moral subjectivism is not a viable position, then we will have found very straightforward ways of defining two sorts of amoral agent. First, the egoist (and the fanatic) is amoral because, whether or not he operates within an objectivist view of values, his concerns are inconsistent with his giving as much consideration to the welfare of others as is required for him to count as a moral agent. Secondly, the rational and reflective subjectivist must be amoral if moral subjectivism is untenable. As I will try to show, moral subjectivism is possible; if not for a very large number of persons in any actual situation, then at least for certain individuals. This leaves the question, which I will attempt to answer in the third section, of how the amoralist is to be distinguished from the moral subjectivist.

Later sections of this chapter deal with some strands of thought in the existentialist tradition which are relevant to my conception of amorality.

2 Forms of moral subjectivism

Is it possible for someone to accept the subjectivist outlook regarding the nature of values and reasons which I outlined in the last chapter and yet be a moral agent? That is, assuming my characterization of the moral agent in Chapter I was acceptable, can someone exhibit all of the qualities of the typical moral agent except the one which requires that he view

his and others' moral judgments as objectively valid or invalid utterances? In particular, can he display an appropriate concern for the welfare of other persons, enter into moral debate, use moral concepts and language in a way which is not grossly misleading or insincere and deliberate from the moral point of view?

Of course, to the extent that moral language is objectivist, and is used by and large by people who suppose themselves to be saying things which can be true or false, or correct or incorrect and so on, the moral subjectivist will be misunderstood and cannot fully commit himself in moral discourse. This, however, need not be a conclusive consideration against his being a moral agent provided the meaning he can give moral concepts is close enough to the ordinary meaning that the misunderstandings are not disruptive of the function of moral discourse and close enough that the moral subjectivist can properly be viewed as participating in the spirit of the moral life.

It is clear, I think, that it will not be very satisfactory to suggest that someone might just happen to value (directly) the welfare of others in the way required to induce him to deliberate from the moral point of view. For one thing, this gives us no account of why such a person would use moral language and in particular of what he might mean by such words as "obligation," "morally right," and "ought." If it is just true that he is concerned to act in ways which take the interests of others into account, we have no explanation of his views (and presumably as a moral agent he will have them) on what

others *ought* to do, and for that matter of his views on what he *ought* to do. Furthermore, unless some fuller account is available, there will be no reason to expect any particular stability in such a person's valuations.

What we need is something which gives the required "moral" valuations of a person a structure and stability which makes sense of the whole range of moral activities in which moral agents engage. One possibility would be to suppose that someone might hold certain deep or fundamental values which he thinks virtually everyone else shares and which underwrite the moral enterprise.

Suppose, for example, that someone values the existence of (and presumably his participation in) a relatively orderly, hospitable and co-operative social environment in which people can exist secure in the expectation that they will not suffer extreme unhappiness or injury at the hands of their fellows. Something like this can reasonably be attributed as a value to the vast majority of people and the fact that realization of the value depends on the effort of virtually everyone, makes it a likely candidate for the present task.

Most people would probably suppose that not only do they value the existence of a cooperative social order but that it is valuable objectively either in itself or because it leads to human well-being or allows people to achieve their highest human potential. That is, most people do suppose that morality rests on or derives from the value of a certain sort of human existence. The moral subjectivist believes that it is only in

the fact that people actually *do* value things that their value resides, since he rejects the notion that things could be valuable in any other "objective" way. But this difference need not be crucial.

A "deep value" explanation of this sort may go some way toward explaining someone's displaying a concern for others, his deliberating from the moral point of view, and his use of moral language. But his engagement in morality must be premised not only on his own deep values and not only on his belief that others share those deep values but on something else as well. There must be an appropriate connection between these deep values and moral thinking. While it may be true that everyone's thinking (and acting) morally is the only means to the achievement of the deep values, the means/end connection will not hold in particular cases for the most part. The connection, from the point of view of the individual, then, cannot simply be that moral considerations have a hold on action because they bear directly on the attainment of some universally valued state of affairs.

The deep values will certainly not be anyone's only values and cases where morality (supported, as it were, by the deep values) requires one action and another value (say personal happiness) requires another will be fairly common. Because the means/end relationship between morally required action and the attainment of a deep value is so tenuous, there is a problem about how moral considerations can come to be very powerful. What is required is that an individual must reason in some way

other than the means/end one. He must take facts about what is required of most people (which ones unspecified) for the deep values to be realized, and treat them as reasons *for him* to do those things. He must ask himself, roughly, "What if everyone (or most people) did that?" and take the answer as being directly relevant to what he is to do.

There is nothing in the subjectivist account of reasons and values I sketched in the last chapter which rules this sort of reasoning out. Indeed it was entirely general as regards what someone might count as a reason for anything. I think that there is, in the form of reasoning required for someone to be a moral subjectivist, an important clue to the best way to understand him. The notion of a deep value explanation is not, I think, very helpful, in itself, precisely because it fails to account for the mode of reasoning required to give the deep value a practical foothold.

The clue is that the reasoning from what is required of people in general to what to do as an individual makes sense only if the deliberating agent has a particular sort of self-conception. He must see himself as essentially a *member* of the whole. That is he must identify himself, to a significant extent, with the whole (the group, society, mankind or whatever) so that whatever reasons apply to the whole apply to him as a member of the whole. This is the only way I can see of closing the gap which exists in any deep value explanation of moral subjectivism.

There are obviously many ways of seeing oneself as a member

but I have in mind one particular way and it involves the notion of one's concept of self. In an interesting article "Ethical Egoism's Brief and Mistaken History," Erling Skorpen (1969) argues that "the concept of the ego or self is non-existent until well over a century after Machiavelli" (pp. 448-9). Up until that time, Skorpen claims, people tended strongly to think of themselves in essentially social terms as members of their respective families, tribes, societies, and so on. If this is true and people tended to deliberate in terms of what "we" ought to do to achieve some end of "ours," the virtual absence of philosophical debate on the possibility of amorality and even on egoism until recently is not surprising.

The conception of self which strongly assimilates the practical "I" and the practical "we" is inherently moral in the sense that, provided only that the "we" is broad enough, that self-conception virtually assures that deliberation will be moral in nature. The moral subjectivist who sees himself and others as members of an organic social unity need not suppose either that this way of seeing things is the only one or the correct one, although in order to engage in moral discourse with others, he must suppose they think in similar terms. And at least in the context of moral discussion others will by and large think in just such terms because of the social, public nature of moral problems. The subjectivist's subjectivism will hardly be noticable since, while it is true that it is only his own valuation of the ends of the social unity which ground their value for him, those ends gain stability and a non-personal

appearance by being perceived through a special self-conception which identifies the self with a group which shares those ends.

Subjectivism, then, seems to be consistent with a stable moral outlook but so far as I can see it requires the adoption of a particular self-conception. There remains, however, the question of how a moral subjectivist can use the notions of obligation and duty in the absence of the concept of objective, interpersonally valid values. I have suggested that anyone who is comfortable with the idea that there are objective values can refer to those values as grounds for his own use of those concepts. Moore, for example, could say:

Our 'duty,' therefore, can only be defined as that action, which will cause more good to exist in the Universe than any possible alternative. And what is 'right' or 'morally permissible' only differs from this, as what will *not* cause *less* good than any possible alternative. When, therefore, Ethics presumes to assert that certain ways of acting are 'duties' it presumes to assert that to act in those ways will always produce the greatest possible sum of good.

(Moore, 1903, *Principia Ethica*, p. 148)

One attempt to make sense of moral obligations in the absence of an acceptance of objective values is found in the development of the concept of institutional obligation. (See especially the Hare/Searle controversy: Hudson (1969).) Basically the idea is that obligations (or many of them) arise because of the existence in society of conventional or rule-governed institutions. To promise, for example, is to invoke the rules of the promise institution, which rules include the requirement that one do as one promised to do. Promising puts one under an obligation by

bringing the rules of the promise institution to bear.

The main line of debate over the concept of institutional obligation concerns the question of whether these obligations are moral or non-moral (or more generally, whether statements about obligations which have been incurred institutionally, are evaluative or descriptive). The debate is complex, but I think it is clear that the (moral) judgment that one ought to keep promises or that someone ought to keep a particular promise, cannot be translated properly into any statement *about* what the promise institution requires. Rather one must view such a statement as one made from *within the institution*.

There are two ways to go from here. One could go on to argue that a judgment made from within an institution is evaluative (or moral) since making the judgment implies a belief that the institution (and hence its rules) are morally justified or required. Another possibility would be to concentrate on the self-conception of someone who makes judgments from within an institution. A subjectivist moral agent might see himself as part of (as involved in) the institution whose rules he is invoking rather than as endorsing it, as it were, from without. If he sees others as similarly part of the institution, he can see his judgments about obligation as invocations of the rules by which all are bound. As long as his option to leave the institution remains in the background he can speak as though the obligations are objectively required. I link his involvement in the institution with his self-conception because as soon as he begins to see himself as an individual who may choose to opt

out of the institution, the obligations lose their objective appearance. As long as he thinks of himself as part of the institution, however, his objectively flavoured judgments will seem quite natural.

The moral subjectivist can, then, participate in moral discourse without grossly misleading other moral agents if he maintains a certain conception of himself as an integral part of society and its institutions.

3 Moral and amoral subjectivism contrasted

It has turned out that subjectivism is not strictly inconsistent with moral agency, but our discussion of moral subjectivism suggests a way of distinguishing the amoral from the moral subjectivist. Where the moral agent sees himself as essentially a part or member of some fairly large social unit, the amoralist sees himself as essentially an individual. Whatever the amoralist's relationship with other persons, he retains a strong sense of his unique individuality. He does not deliberate in the first person plural mode nor by asking questions like "What if everyone does this?", because that would not solve his first person practical problems. Where the moral agent sees the laws, institutions and customs of his society as part of his society's (and hence his) attempt to achieve certain ends, the amoralist sees these as external constraints and artifacts around which or through which he must plan his own projects. The amoralist is a member of society in many ways but in spirit he remains an outsider. Even when his projects coincide with those of others

or even of society as a whole he never quite loses sight of them as *his* projects, first and foremost.

There are interesting, conceptual links among the ideas of individualism, subjectivism and amoralism, but I will not attempt to trace them all here. It is no accident that the rise of individualistic political philosophy in the western world has been accompanied by moral skepticism and a movement away from traditional moral values (at least a questioning of the latter). It is interesting as well that much of the reaction against the prevailing moral climate has taken the form of an idealization of just the kind of self-conception I have been imputing to the moral subjectivist, as if in recognition of the fact that the only way to stop the shift to amoralism prompted by a general sympathy with (although I think not a genuine acceptance of) subjectivism in values, is to push the objectivism/subjectivism issue into the background by changing the mode in which practical problems are perceived.

The amoralist is a conceptual individualist. He sees himself as standing apart from others and as creating and sustaining his value framework and because of this his position is inherently unstable. There is always a pressure on the deviant to conform and it is only through a sustained effort and continual reflection that the self-conception of the amoral agent can be sustained. This instability is, however, shared by the moral subjectivist because of what I have referred to as the natural drift from moral thinking to objectivism.

4 Existentialism

There are two main streams of modern philosophical thought which have taken very seriously the idea of the subjectivity of values; emotivism and existentialism. Yet in spite of this common concern, there could hardly be found two schools of thought with more divergent styles and methods. Ethical emotivists following the lead of Hume, and later of Ayer and Stevenson typically argue that their theories are essentially attempts to describe ordinary moral discourse and have generally supposed that nothing in their theories ought to upset the ordinary man or make him change his attitudes and outlook. The spirit of existentialism, on the other hand, is much more activist and much of the existentialist literature is aimed at a broader readership than just academic philosophers. The existentialist wants to understand "the phenomena" and not the language (which may or may not reflect the true nature of things) and the understanding gained may, he allows, produce significant changes in the outlook of the ordinary man.

I have already said something about emotivist theories and I have tried to deal with them pretty much on their own terms, treating them as descriptive analyses of ordinary moral discourse. I should now like to say a little about the existentialist approach to subjectivism. Actually, of course, existentialist thinkers do not form a very coherent group and it would be a mistake to think that very much could be said which would hold for all of them. If there is a crucial centre to existentialist thought it is the notion of individual freedom of choice at a

very fundamental level and along with it a belief that people must make such fundamental choices for themselves. They are not to be made by any outside or objective considerations. However, the result is not always a Sartrian subjectivism, especially when the existentialism is religious. I quote a passage from Martin Buber:

"If I have done away with God the Father....," Sartre says literally, "someone is needed to invent values.... Life has no meaning a priori ... it is up to you to give it a meaning, and value is nothing else than this meaning which you choose." That is almost exactly what Nietzsche said, and it has not become any truer since then. One can believe in and accept a meaning or value, one can set it as a guiding light over one's life if one has discovered it, not if one has invented it. It can be for me an illuminating meaning, a direction-giving value only if it has been revealed to me in my meeting with Being, not if I have freely chosen it for myself from among the existing possibilities.

(Buber, 1957, *Eclipse of God*, pp. 68-70)

Thus, the existential emphasis on the subjectivity of the individual and even on the necessity of fundamental choice does not lead necessarily to a subjectivist view of values. It may be that the fundamental choice is seen to be the "unreasoned" leap into faith from which point objective value can be discovered.

In any case, the diversity among existentialists being what it is, it is best to deal with just one writer and since J.P. Sartre is probably the most widely read and is squarely in the main stream of existentialist thought, I propose to concentrate on his contribution for the most part. Sartre's ethical theory in *Being and Nothingness* (if it can be called an ethical theory) grows out of his analysis of the nature of consciousness. I

cannot hope here to present Sartre's theory and I shall have to assume that it is familiar enough for my purposes. The main question I wish to address is whether or not Sartre's views, and insofar as his views are typical of existentialists, whether or not existentialist theories in general, are consistent with the notion of (subjectivist) moral agency.

5 Jean-Paul Sartre

It is clear, I think that Sartre's brand of existentialism, at least as it is expounded in *Being and Nothingness*, is not a moral theory, nor even a metamoral theory. That it is not a moral theory Sartre himself recognizes, and in the last section of the book he points out that the only implication of ontology and existential psychoanalysis (which he calls "moral description") is to reveal to people the ultimate subjectivity of values. Sartre seems to take it for granted that knowing the truth about these matters will have some effect on how we think about what we do, whether or not it makes any difference to what we do. But even if the acceptance of his theory does affect what people do (perhaps by freeing them from bad faith and a belief in objective value, and thus opening up a new range of choices where none were recognized before) those consequences would not be intended by or implied in the theory itself and so would not suffice to constitute the theory as an ethical or moral one.

Yet, while Sartre's theory is normatively silent on various sorts of interpersonal relations it is not altogether silent descriptively. He discusses love, hate, lust, desire and so on,

often in great depth but, except for the occasional mention of bad faith, the normative issues remain in the background. Indeed, with the exception of the implicit demand that we see things as they are (as Sartre claims they are), there are not even any formal requirements on what we can value or what projects we can undertake. Our freedom is infinite although bounded by facticity. It is interesting that in his *Existentialism* (which he later repudiated), Sartre does advance a formal constraint on valuation and he does so in order to answer the charge that existentialism leaves man free to choose his values irresponsibly and arbitrarily without regard for human solidarity. I have already quoted two of the relevant passages earlier and it will be recalled that the claim in these was that, in some cases, in choosing a value or an ideal, a man chooses it as good universally, as good for all, as an ideal for man as such, and so on. This maneuver, however, is not only contradictory to the spirit of much of *Being and Nothingness*, but seems to be quite unfounded on anything at all in the rest of Sartre's theory. It looks very much like an attempt to show that his existentialism is at least a quasi-moral theory after all. It is as if Sartre realized that unless something of this sort could be established, the adoption of an existentialist outlook would tend toward amorality.

I suggested earlier in this chapter that the most promising form of subjectivist moral agency was the one which involved the individual in a particular self-conception, e.g. as part of a collective "we." Interestingly, in spite of the fact that Sartre's theory in *Being and Nothingness* pushes in exactly the opposite

direction, he ended up in *The Critique of Dialectical Reason* defending a Marxist position in which this possibility of seeing oneself as a *member*, as a part of a group with a coincidence of interests, is absolutely critical.

In *Being and Nothingness*, Sartre (1966) says:

The spirit of seriousness has two characteristics: it considers values as transcendent givens independent of human subjectivity, and it transfers the quality of "desirable" from the ontological structure of things to their simple material constitution. For the spirit of seriousness, for example, *bread* is desirable because it is *necessary* to live (a value written in an intelligible heaven) and because bread *is* nourishing. The result of the serious attitude, which as we know rules the world, is to cause the symbolic values of things to be drunk in by their empirical idiosyncrasy as ink by a blotter; it puts forward the opacity of the desired object and posits it in itself as a desirable irreducible. Thus we are already on the moral plane but concurrently on that of bad faith, for it is an ethics which is ashamed of itself and does not dare speak its name. It has obscured all its goals in order to free itself from anguish (p. 796).

It almost seems that for Sartre the "moral plane" is both objectivist and necessarily in bad faith.

Mary Warnock in her excellent book *Existentialism* discusses the question of whether or not the theory of *Being and Nothingness* is consistent with moral agency. She says:

There is a real difficulty at this point in Sartre's philosophy which he did not in any way solve, at least until he abandoned Existentialism. In *Being and Nothingness* he seems to be saying that we must each decide for ourselves how to live, what is good and what is bad, and that this is a purely personal decision, which no one can take on behalf of another. But there is an element in genuine evaluation which will

not submit to this analysis. If, for example, a man judges sincerely that tax evasion is wrong, then in some sense he has, whether he knows it or not, judged that it is wrong in general, and he may even believe, though without saying as much, that it is wrong necessarily, or absolutely.... And so Sartre has not said enough, when he insists that human beings cannot find absolute values in the world, they can only pretend to themselves to do so. He has not taken into account of the facts of forming moral opinions.

(Warnock, 1970, p. 124)

I think that Warnock here voices a very common reaction to Sartre's writing insofar as she is unable to see even the possibility of moral thinking within Sartre's terms of reference. I have used the terms "valuation" throughout instead of "evaluation" to mark just the sort of difference which she points up between the existentialist (amoral) approach and the moral approach to values.

Sartre, himself has a very active social conscience and is politically very aware. If Warnock is right (and I agree with her here) that Sartre's existentialism is essentially amoral (my term, not hers), then it is not surprising that he should find it difficult to continue to embrace a philosophy which makes no special provision for, and is even possibly inconsistent with, the sincere use of moral language. Moral language is after all the language of socially concerned political man. I should say that I am not suggesting that the amoralist must not engage in the affairs of society but only that he must do so with one eye firmly fixed on himself as an individual existent who in Sartre's phrase is "the being by whom values exist."

6 Nietzsche

A natural place to look for a treatment of the possibility of abandoning morality is in the writings of Friedrich Nietzsche. Unfortunately it is not easy, having looked, to say what is there. Certainly there is an important theme of anti-moralism in Nietzsche's thought and while some commentators have construed him as advocating a new and different morality, it seems to me that his rejection of conventional moral values is so radical and thorough that it would be strained to speak of morality at all regarding his alternative view.

The concept on which we must concentrate is clearly that of the "ubermensch" or "overman" but first some remarks about Nietzsche's hostility toward conventional morality are in order since, in large measure, the overman is the embodiment of what Nietzsche sees as the way of going beyond ordinary moral thinking. Here, as elsewhere, the issue is clouded by Nietzsche's tendency to theorize in terms of a biological psychology. He says, for example, that what people value is determined by their psychological makeup; indeed that the morality of a group could no more be changed than could its physique or the biochemical composition of its blood. The whole doctrine of the will to power is saturated with this sort of naturalism. However, it is possible to ignore most of Nietzsche's psychological hypotheses while preserving a reasonably coherent point of view, because, in spite of them, he certainly thought that it is possible to conduct a critique of morality and the results of the critique might be expected to influence our ways of thinking and acting.

In the Preface to the *Genealogy of Morals* (6), he says:

Let us articulate this *new demand*: we need a *critique of moral values, the value of these values themselves must first be called in question....* One has taken the value of these "values" as given, as factual, as beyond all question; one has hitherto never doubted or hesitated in the slightest degree in supposing "the good man" to be of greater value than "the evil man," of greater value in the sense of furthering the advancement and prosperity of man in general (the future of man included). But what if the reverse were true?.... So that precisely morality would be to blame, if the *highest power and splendor* actually possible to the type man was never in fact attained? So that precisely morality was the danger of dangers.

(Nietzsche, 1969, p. 20)

Notice the shift from "prosperity" to "power and splendor" as well as the shift from "man in general" to "the type man."

Nietzsche often denegrates the typically moral concern with comfortable living, security, happiness, prosperity, and so on, in favour of values which are better thought of as human excellences than moral virtues; e.g. self-discipline, intellectual conscientiousness, self-sufficiency, and so on. Consider the following passage:

You want, if possible - and there is no more insane "if possible" - *to abolish suffering*. And we? It really seems that we would rather have it higher and worse than ever. Well-being as you understand it - that is no goal, that seems to us an *end*, a state that soon makes man ridiculous and contemptible - that makes his destruction *desirable*.

(Nietzsche, 1966, *Beyond Good and Evil*, 225, p. 153).

I think there can be little doubt that Nietzsche, the self-confessed amoralist, is prepared to take very seriously indeed

the possibility of abandoning morality itself, and not just of modifying conventional morality.

It is interesting, however, that he does not envisage any real possibility of a universal abandonment of morality because the "herd" can hardly be expected to exemplify nor even appreciate human excellence (as he understands it) to any great extent. Sometimes he even suggests that morality *ought* to prevail in the herd as the condition of the existence of the overman. He says:

The continuance of the Christian ideal is one of the most desirable things there are - even for the sake of the ideals that want to stand beside it and perhaps above it - they must have opponents, strong opponents, if they are to become *strong*.

Thus we immoralists require the power of morality: our drive of self-preservation wants our *opponents* to retain their strength - it only wants to become *master over them*.

(Nietzsche, 1968, *The Will to Power*, 361, p. 197)

At this point, we should note a tension, if not a contradiction, in Nietzsche's thinking. On the one hand, he resists the notion that there are objective values to be discovered by the contemplation of the nature of things, as it were. Not only must the moral values be revalued, but all values must constantly be revalued. For Nietzsche, much more than for Sartre, this "must" is ethical and thoroughly normative, since it is only by adopting this ever-questioning point of view that a man can achieve his highest potential. Nietzsche has, it seems, a view of the ideal, highest, best kind of existence for man and it is an ideal which, while not attainable for any but a small elite, is nonetheless very much like an objective value of the sort he

denies exists. Indeed, Nietzsche claims that the existence of one overman could justify the world! He says:

I teach: that there are higher and lower men, and that a single individual can under certain circumstances justify the existence of whole millennia - that is, a full, rich, great, whole human being in relation to countless incomplete fragmentary men.

(Nietzsche, 1968, *The Will to Power*, 997, p. 518)

It is comments such as these which have convinced some people that Nietzsche is actually a moralist in spite of himself. It seems to me better to say that he may be an objectivist or even that he has an ethical theory, but that the nature of his ethics does not permit the title "moralist." (See *The Will to Power*, 304, p. 171 where Nietzsche suggests that even moralists, i.e. those who would propagate moral virtues, actually need to be amoralists so as not to succumb to their own virtues and the notion that they embody objective values.) If nothing else, his lack of concern and the overman's lack of concern, for the well-being of the bulk of mankind disqualifies him on that score. But, in any case, it is not Nietzsche we are primarily concerned with here; it is the overman, and surely he does not even begin to qualify as a moral agent.

The overman stands alone. Again and again Nietzsche paints him as the solitary individual, as for example in the following passages:

I teach: the herd seeks to preserve one type and defends itself on both sides, against those who have degenerated from it (criminals, etc.) and those who tower above it. The tendency of the herd is directed toward standstill and preservation, there

is nothing creative in it.

The pleasant feelings with which the good, benevolent, just man inspires (in contrast to the tension, fear which the great, new man arouses) are our own feelings of personal security and equality: the herd animal thus glorifies the herd nature and then it feels comfortable. This judgment of comfort masks itself with fair words - thus "morality" arises. - But observe the hatred of the herd for the truthful. - (285, p. 162).

Let one not be deceived about oneself! If one hears within oneself the moral imperative as it is understood by altruism, one belongs to the herd. If one has the opposite feeling, if one feels one's danger and aberration lies in disinterested and selfless actions, one does not belong to the herd. (286, p. 162)

(Nietzsche, 1968, *The Will to Power*)

Thus the overman does not display the kind of consciousness of community which is one hallmark of the moral agent. Further, as I have noted already, the overman does not accept the objectivity of values; he never loses sight of the fact that his values are the result of *his* valuations and as such they must constantly be affirmed in action - otherwise they lose their vitality and "degenerate into truth."

Finally, the overman does not regard others, or their suffering as having any claims on his actions. It is no use against this contention to quote Nietzsche's occasional remarks about the magnanimity and generosity of the higher type. If the overman benefits mankind, it is only because he fulfills himself in this way. Nietzsche even suggests that it is the will to power which prompts the overman to sacrifice because it takes most strength to do what does not come naturally and it takes strength to

squander one's resources - especially resources of the spirit. Whatever the "duty" the exceptional man has toward lesser men, it is not a moral duty, being closer to *noblesse oblige* than to a respect for humanity in other persons.

There is much in Nietzsche's discussions of the overman which I do not understand, e.g. the role of passion versus reason, and the operation of instinct. However, I think it is clear that the overman is indeed an amoralist and whether we find ourselves drawn toward him or repulsed, it must be admitted that the possibility of someone's abandoning morality in this way is an interesting one. I should say that the overman is only one sort of amoral person - we have encountered others earlier - and he is in fact the result of Nietzsche's view of the ideal human type. Not everyone will share this ideal, of course, but the main interest for me in Nietzsche's conception of *das Uebermensch* is that it points out quite clearly that the notion of an amoral agent can be filled out in a way which can capture the imagination of at least one great thinker. I suspect it is capable of inspiring many other people as well, including some who could not quite admit it. Morality is not the only interesting and rational human possibility.

BIBLIOGRAPHY

- Aiken, Henry David. "The Concept of Moral Objectivity." Morality and the Language of Conduct, Hector-Neri Castañeda and George Nakhnikian, eds. Detroit, Wayne State University Press, 1963, pp. 69-105. (Reprinted from H.D. Aiken. Reason and Conduct. N.Y., Knopf, 1962.)
- Allport, G.W. "Normative Compatibility in the Light of Social Science." New Knowledge in Human Values, A.H. Maslow, ed. N.Y., Harper and Bros., 1959, pp. 137-150.
- Ayer, A.J. "On the Analysis of Moral Judgements." Philosophical Essays, London, Macmillan; N.Y., St. Martin's Press, 1963, pp. 231-249 (from Horizon, vol. 20 (1949))
- Baier, Kurt. The Moral Point of View: A Rational Basis of Ethics. Ithaca, N.Y., Cornell University Press, 1958.
- _____. "Ethical Egoism and Interpersonal Compatibility." Philosophical Studies, vol. 24 (1973), pp. 357-368.
- Barrett, William. Irrational Man. Garden City, N.Y., Doubleday Anchor Books, 1962 (copyright, 1958).
- Baumer, William H. "Indefensible Impersonal Egoism." Philosophical Studies, vol. 18 (1967), pp. 72-75.
- Beehler, Rodger. "Morals and Reasons." Analysis, vol. 33 (1972-1973), pp. 19-21.
- _____. "Reasons for Being Moral." Analysis, vol. 33 (1972-1973), pp. 12-16.
- Berkeley, George. Three Dialogues Between Hylas and Philonous. in Berkeley's Philosophical Writings, D.M. Armstrong, ed. London, Collier-Macmillan, 1965, pp. 129-225.
- Branden, Nathaniel. "Rational Egoism." The Personalist, vol. 51 (1970), pp. 305-314.
- Brentano, Franz Clemens. The Origin of our Knowledge of Right and Wrong. Oskar Kraus, ed. English edition, Roderick Chisholm, ed. Roderick Chisholm and Elizabeth H. Schneewind, eds. London, Routledge and K. Paul; N.Y., Humanities Press, 1969.

- Broad, Charles Dunbar. Five Types of Ethical Theory. London, K. Paul, Trench, Trubner and Co.; New York, Harcourt, Brace and Co., 1930.
- Buber, Martin. Eclipse of God. Studies in the Relation between Religion and Philosophy. Maurice Friedman et al., trans. New York, Harper Torchbooks, 1957.
- Butler, Joseph. Fifteen Sermons. W.R. Matthews ed. London, G. Bell and Sons, 1967.
- Campbell, Richmond. "A Short Refutation of Ethical Egoism." Canadian Journal of Philosophy, vol. 2 (1972) pp. 249-254.
- Clarke, Samuel. "Discourse upon Natural Religion." British Moralists, vol. 2, L.A. Selby-Bigge, ed. N.Y., Dover Publications, 1965, pp. 482-525.
- Cleckley, Hervey. The Mask of Sanity. 5th ed. St. Louis, C.V. Mosby Co., 1976.
- Edwards, Paul. The Logic of Moral Discourse. Introduction by Sidney Hook. N.Y., The Free Press; London, Collier-Macmillan, 1955.
- Emmons, Donald. "Refuting the Egoist." The Personalist, vol. 50 (1969), pp. 309-319.
- Falk, W.D. "'Ought' and Motivation." Proceedings of the Aristotelian Society, vol. 48 (1947-8), pp. 111-138.
- _____. "Morality, Self and Others." Morality and the Language of Conduct, Hector-Neri Castañeda and George Nakhnikian, eds. Detroit, Wayne State University Press, 1963, pp. 25-67.
- Firth, Roderick. "Ethical Absolutism and the Ideal Observer." Philosophy and Phenomenological Research, vol. 12 (1951/1952), pp. 317-345.
- Foot, Philippa. "Moral Beliefs." Proceedings of the Aristotelian Society, vol. 59, (1958-9), pp. 83-104.
- Gauthier, David. "Morality and Advantage." Philosophical Review, vol. 76 (1967), pp. 460-475.
- _____. "The Impossibility of Rational Egoism." The Journal of Philosophy, vol. 71 (1974), pp. 439-456.
- _____. "Reason and Maximization." Canadian Journal of Philosophy, vol. 4 (1974-5), pp. 411-433.
- _____. "Coordination." Dialogue, vol. 14 (1975), pp. 195-221.

- Gert, Bernard. The Moral Rules. New York, Harper and Row, 1966.
- Glasgow, W.D. "The Contradiction in Ethical Egoism." Philosophical Studies, vol. 19 (1968), pp. 81-85.
- Hall, Everett Wesley. What Is Value? An Essay in Philosophical Analysis. London, Routledge and Paul, 1952.
- _____. Our Knowledge of Fact and Value. Introduction by E.M. Adams. Chapel Hill, University of North Carolina Press, 1961.
- Hare, R.M. The Language of Morals. Oxford, Oxford University Press, 1952. (Oxford University Press paperback, 1964).
- _____. Freedom and Reason. Oxford, Oxford University Press, 1963. (Oxford University Press paperback, 1965).
- Harman, Gilbert. The Nature of Morality. New York, Oxford University Press, 1977.
- Hobbes, Thomas. Leviathan. Introduction by A.D. Lindsay. London, Dent; New York, Dutton, Everyman's Library (691), 1914.
- Hospers, John. "Baier and Medlin on Ethical Egoism." Philosophical Studies, vol. 12 (1961), pp. 10-16.
- Hudson, W.D. (ed.) The Is/Ought Question. London, Macmillan, 1969.
- Hume, David. A Treatise of Human Nature. L.A. Selby-Bigge, ed. Oxford, Clarendon Press, 1888.
- _____. An Enquiry Concerning the Principles of Morals. 2nd edition, (reprinted from the edition of 1777), LaSalle, Illinois, Open Court, 1966.
- Jack, Henry. "The Consistency of Ethical Egoism." Dialogue, vol. 8 (1969), pp. 475-480.
- Kant, Immanuel. Groundwork of the Metaphysic of Morals. H.J. Paton, trans. New York, Harper and Row (Harper Torchbooks), 1964.
- Kierkegaard, Søren. Fear and Trembling. Walter Lowrie, trans. with introduction and notes by W. Lowrie. Princeton, Princeton University Press, 1941.
- Lewis, Clarence Irving. The Ground and Nature of the Right. New York, Columbia University Press, 1955.
- _____. Values and Imperatives: Studies in Ethics. John Lange, ed. Stanford, California, Stanford University Press, 1969.

- Mackie, J.L. Ethics: Inventing Right and Wrong. Penguin Books, 1977.
- Medlin, Brian. "Ultimate Principles and Ethical Egoism." Australian Journal of Philosophy, vol. 35 (1957), pp. 111-118.
- Moore, George Edward. Principia Ethica. Cambridge, Cambridge University Press, 1903.
- Nagel, Thomas. The Possibility of Altruism. Oxford, Clarendon Press, 1970.
- Narveson, Jan. Morality and Utility. Baltimore, John Hopkins Press, 1967.
- Nielsen, Kai. "Is 'Why Should I be Moral?' an Absurdity?" Australasian Journal of Philosophy, vol. 36 (1958), pp. 25-32.
- _____. "Why Should I be Moral?" Methodos, vol. 15 (1963), pp. 275-306.
- _____. "On Giving Reasons for Being Moral." Analysis, vol. 33 (1972-1973), pp. 17-19.
- Nietzsche, Friedrich. Beyond Good and Evil. trans. with commentary by Walter Kaufmann. N.Y., Random House, Vintage Books, 1966.
- _____. The Will to Power. W. Kaufmann and R.J. Hollingdale, trans., ed. with commentary by Walter Kaufmann. N.Y., Random House, Vintage Books, 1968 (copyright, 1967).
- _____. On the Genealogy of Morals and Ecce Homo. W. Kaufmann and R.J. Hollingdale, trans., Walter Kaufmann, ed. N.Y., Random House, Vintage Books, 1969 (copyright, 1967).
- _____. Thus Spoke Zarathustra. trans. with an introduction by R.J. Hollingdale. Penguin Books, 1969.
- Paton, H.J. "The Emotive Theory of Ethics," Symposium. Logical Positivism and Ethics. Aristotelian Society, Supplementary Volume 22. London, Harrison and Sons, 1948, pp. 107-126.
- Plato. The Republic of Plato. F.M. Cornford, trans. Oxford, Oxford University Press, 1941. (Oxford University Press paperback, 1945).
- Prichard, H.A. "Does Moral Philosophy Rest on a Mistake?" Mind, vol. 21 (1912), pp. 21-37.

- Prichard, H.A. Moral Obligation. Oxford, Clarendon Press, 1949.
- Prior, Arthur N. Logic and the Basis of Ethics. Oxford, Clarendon Press, 1949.
- Quinn, Warren. "Egoism as an Ethical System." The Journal of Philosophy, vol. 71 (1974), pp. 456-472.
- Rand, Ayn. The Virtue of Selfishness. New York, New American Library (Signet), 1964.
- Robinson, Richard. "The Emotive Theory of Ethics," Symposium. Logical Positivism and Ethics. Aristotelian Society, Supplementary Volume 22. London, Harrison and Sons, 1948, pp. 79-106.
- Sartre, Jean-Paul. Existentialism. Bernard Frechtman, trans. New York, Philosophical Library, 1947.
- _____. Being and Nothingness. trans. and with an introduction by Hazel E. Barnes. N.Y., Washington Square Press, Pocket Books, 1966.
- _____. Critique of Dialectical Reason, Theory of Practical Ensembles. Jonathan Rée, ed. London, NLB; Atlantic Highlands, N.J., Humanities Press, 1976.
- Scruton, Roger. "Attitudes, Beliefs and Reasons." Morality and Moral Reasoning, John Casey, ed. London, Methuen and Co., 1971, pp. 25-100.
- Selby-Bigge, L.A. ed. British Moralists. 2 vols. New York, Dover Publications, 1965.
- Sidgwick, Henry. The Methods of Ethics. 7th ed. New York, Dover Publications, 1966.
- Silverstein, Harry S. "Universalizability and Egoism." Australasian Journal of Philosophy, vol. 46 (1968), pp. 242-264.
- Singer, Marcus G. "On Duties to Oneself." Ethics, vol. 69 (1959), pp. 202-205.
- Singer, Peter. "The Triviality of the Debate over 'Is-Ought' and the Definition of 'Moral'." American Philosophical Quarterly, vol. 10 (1973), pp. 51-56.
- Skorpen, Erling. "Ethical Egoism's Brief and Mistaken History." The Personalist, vol. 50 (1969), pp. 448-472.
- Smith, Adam. The Theory of Moral Sentiments. D.D. Raphael and A.L. Macfie eds. Oxford, Clarendon Press, 1976.

- Stevenson, Charles L. Ethics and Language. New Haven and London, Yale University Press, 1944.
- Strawson, P.F. "Social Morality and Individual Ideal." Philosophy, vol. 36 (1961), pp. 1-17.
- Taylor, Richard. Good and Evil: A New Direction. New York, Macmillan, 1970.
- Toulmin, Stephen Edelston. An Examination of the Place of Reason in Ethics. Cambridge, Cambridge University Press, 1968.
- Warnock, Mary. Existentialism. Oxford, Oxford University Press, 1970.
- Wittgenstein, Ludwig. The Blue and Brown Books. N.Y., Harper and Row, Harper Torchbooks, 1965.