

ETHICS

By G. E. MOORE

LONDON

WILLIAMS & NORGATE

HENRY HOLT & Co., NEW YORK

CANADA: WM. BRIGGS, TORONTO

INDIA: R. & T. WASHBOURNE, LTD.



HOME
UNIVERSITY
LIBRARY
OF
MODERN KNOWLEDGE

Editors:

HERBERT FISHER, M.A., F.B.A.

PROF. GILBERT MURRAY, D.LITT.,
LL.D., F.B.A.

PROF. J. ARTHUR THOMSON, M.A.

PROF. WILLIAM T. BREWSTER, M.A.
(Columbia University, U.S.A.)

NEW YORK
HENRY HOLT AND COMPANY



ETHICS

BY

G. E. MOORE

LECTURER IN MORAL SCIENCE IN THE
UNIVERSITY OF CAMBRIDGE

LONDON
WILLIAMS AND NORGATE

भारती-भवन पुस्तकालय

प्रयाग

क्रमिक संख्या

विभाग

10227

170/17.

Printed by
MORRISON & GIBB LIMITED
Edinburgh

CONTENTS

CHAP.		PAGE
I	UTILITARIANISM	7
II	UTILITARIANISM (<i>concluded</i>)	40
III	THE OBJECTIVITY OF MORAL JUDGMENTS	79
IV	THE OBJECTIVITY OF MORAL JUDGMENTS (<i>concluded</i>)	133
V	RESULTS THE TEST OF RIGHT AND WRONG	170
VI	FREE WILL	196
VII	INTRINSIC VALUE	223
	NOTE ON BOOKS	253
	INDEX	255

भारती-भवन पुस्तकालय
प्रयाग

क्रमिक संख्या
विभाग

10227

170/17.

Printed by
MORRISON & GIBB LIMITED
Edinburgh

CONTENTS

CHAP.	PAGE
I UTILITARIANISM	7
II UTILITARIANISM (<i>concluded</i>)	40
III THE OBJECTIVITY OF MORAL JUDGMENTS	79
IV THE OBJECTIVITY OF MORAL JUDGMENTS (<i>concluded</i>)	133
V RESULTS THE TEST OF RIGHT AND WRONG	170
VI FREE WILL	196
VII INTRINSIC VALUE	223
NOTE ON BOOKS	253
INDEX	255

ETHICS

CHAPTER I

UTILITARIANISM

ETHICS is a subject about which there has been and still is an immense amount of difference of opinion, in spite of all the time and labour which have been devoted to the study of it. There are indeed certain matters about which there is not much disagreement. Almost everybody is agreed that certain kinds of actions ought, as a general rule, to be avoided ; and that under certain circumstances, which constantly recur, it is, as a general rule, better to act in certain specified ways rather than in others. There is, moreover, a pretty general agreement, with regard to certain things which happen in the world, that it would be better if they never happened, or, at least, did not happen

7

so often as they do; and with regard to others, that it would be better if they happened more often than they do. But on many questions, even of this kind, there is great diversity of opinion. Actions which some philosophers hold to be generally wrong, others hold to be generally right and occurrences which some hold to be evils, others hold to be goods.

And when we come to more fundamental questions the difference of opinion is even more marked. Ethical philosophers have, in fact, been largely concerned, not with laying down rules to the effect that certain ways of acting are generally or always right, and others generally or always wrong, nor yet with giving lists of things which are good and others which are evil, but with trying to answer more general and fundamental questions such as the following. What, after all, is it that we mean to say of an action when we say that it is right or ought to be done? And what is it that we mean to say of a state of things when we say that it is good or bad? Can we discover any general characteristic, which belongs in

common to absolutely *all* right actions, no matter how different they may be in other respects? and which does not belong to any actions except those which are right? And can we similarly discover any characteristic which belongs in common to absolutely all "good" things, and which does not belong to any thing except what is a good? Or again, can we discover any single reason, applicable to all right actions equally, which is, in every case, *the* reason why an action is right, when it is right? And can we, similarly, discover any reason which is *the* reason why a thing is good, when it is good, and which also gives us the reason why any one thing is better than another, when it is better? Or is there, perhaps, no such single reason in either case? On questions of this sort different philosophers still hold the most diverse opinions. I think it is true that absolutely every answer which has ever been given to them by any one philosopher would be denied to be true by many others. There is, at any rate, no such consensus of opinion among experts about these fundamental ethical

questions, as there is about many fundamental propositions in Mathematics and the Natural Sciences.

Now, it is precisely questions of this sort, about every one of which there are serious differences of opinion, that I wish to discuss in this book. And from the fact that so much difference of opinion exists about them it is natural to infer that they are questions about which it is extremely difficult to discover the truth. This is, I think, really the case. The probability is, that hardly any positive proposition, which can as yet be offered in answer to them, will be strictly and absolutely true. With regard to *negative* propositions, indeed,—propositions to the effect that certain positive answers which have been offered, are false,—the case seems to be different. We are, I think, justified in being much more certain that some of the positive suggestions which have been made are *not* true, than that any particular one among them *is* true; though even here, perhaps, we are not justified in being *absolutely* certain.

But ^g if we cannot be justified either in general _c

accepting or rejecting, with absolute certainty, any of the alternative hypotheses which can be suggested, it is, I think, well worth while to consider carefully the most important among these rival hypotheses. To realise and distinguish clearly from one another the most important of the different views which may be held about these matters is well worth doing, even if we ought to admit that the best of them has no more than a certain amount of probability in its favour, and that the worst have just a possibility of being true. This, therefore, is what I shall try to do. I shall try to state and distinguish clearly from one another what seem to me to be the most important of the different views which may be held upon a few of the most fundamental ethical questions. Some of these views seem to me to be much nearer the truth than others, and I shall try to indicate which these are. But even where it seems pretty certain that some one view is erroneous, and that another comes, at least, rather nearer to the truth, it is very difficult to be sure that the latter is strictly and absolutely true.

One great difficulty which arises in ethical discussions is the difficulty of getting quite clear as to exactly what question it is that we want to answer. And in order to minimize this difficulty, I propose to begin, in these first two chapters, by stating one particular theory, which seems to me to be peculiarly simple and easy to understand. It is a theory which, so far as I can see, comes very near to the truth in some respects, but is quite false in others. And why I propose to begin with it is merely because I think it brings out particularly clearly the difference between several quite distinct questions, which are liable to be confused with one another. If, after stating this theory, we then go on to consider the most important objections which might be urged against it, for various reasons, we shall, I think, pretty well cover the main topics of ethical discussion, so far as fundamental principles are concerned.

This theory starts from the familiar fact that we all very often seem to have a choice between several different actions, any one

of which we might do, if we chose. Whether, in such cases, we really do have a choice, in the sense that we ever really *could* choose any other action than the one which in the end we do choose, is a question upon which it does not pronounce and which will have to be considered later on. All that the theory assumes is that, in many cases, there certainly are a considerable number of different actions, any one of which we could do, *if* we chose, and between which, therefore, in *this* sense, we have a choice ; while there are others which we could not do, even if we did choose to do them. It assumes, that is to say, that in many cases, *if* we had chosen differently, we should have acted differently ; and this seems to be an unquestionable fact, which must be admitted, even if we hold that it is never the case that we *could* have chosen differently. Our theory assumes, then, that many of our actions are under the control of our wills, in the sense that *if*, just before we began to do them, we had chosen not to do them, we *should* not have done them ; and I propose to call all actions of this kind *voluntary* actions.

It should be noticed that, if we define voluntary actions in this way, it is by no means certain that all or nearly all voluntary actions are actually themselves chosen or willed. It seems highly probable that an immense number of the actions which we do, and which we *could* have avoided, *if* we had chosen to avoid them, were not themselves willed at all. It is only true of them that they are "voluntary" in the sense that a particular act of will, just before their occurrence, would have been sufficient to *prevent* them; not in the sense that they themselves were brought about by being willed. And perhaps there is some departure from common usage in calling all such acts "voluntary." I do not think, however, that it is in accordance with common usage to restrict the name "voluntary" to actions which are quite certainly actually willed. And the class of actions to which I propose to give the name—all those, namely, which we could have prevented, *if*, immediately beforehand, we had willed to do so—do, I think, certainly require to be distinguished by some special name. It might, perhaps, be thought that

almost all our actions, or even, in a sense, *absolutely* all those, which properly deserve to be called "ours," are "voluntary" in this sense: so that the use of this special name is unnecessary: we might, instead, talk simply of "our actions." And it is, I think, true that almost all the actions, of which we should generally think, when we talk of "our actions," are of this nature; and even that, in some contexts, when we talk of "human actions," we do refer exclusively to actions of this sort. But in other contexts such a way of speaking would be misleading. It is quite certain that both our bodies and our minds constantly do things, which we certainly could not have prevented, by merely willing just beforehand that they should not be done; and some, at least, of these things, which our bodies and minds do, would in certain contexts be called actions of ours. There would therefore be some risk of confusion if we were to speak of "human actions" generally, when we mean only actions which are "voluntary" in the sense I have defined. It is better, therefore, to give some special name to

actions of this class ; and I cannot think of any better name than that of "voluntary" actions. If we require further to distinguish from among them, those which are also voluntary in the sense that we definitely willed to do them, we can do so by calling these "willed" actions.

Our theory holds, then, that a great many of our actions are voluntary in the sense that we could have avoided them, *if*, just beforehand, we had chosen to do so. It does not pretend to decide whether we *could* have thus chosen to avoid them ; it only says that, *if* we had so chosen, we should have succeeded. And its first concern is to lay down some absolutely universal rules as to the conditions under which actions of this kind are *right* or *wrong* ; under which they *ought* or *ought not* to be done ; and under which it is our *duty* to do them or not to do them. It is quite certain that we do hold that many voluntary actions are right and others wrong ; that many ought to have been done, and others ought not to have been done ; and that it was the agent's duty to do some of them, and his duty not to

do others. Whether any actions, except voluntary ones, can be properly said to be right or wrong, or to be actions which ought or ought not to have been done, and, if so, in what sense and under what conditions, is again a question which our theory does not presume to answer. It only assumes that these things *can* be properly said of some voluntary actions, whether or not they can also be said of other actions as well. It confines itself, therefore, strictly to voluntary actions; and with regard to these it asks the following questions. Can we discover any characteristic, over and above the mere fact that they *are* right, which belongs to absolutely *all* voluntary actions which are right, and which at the same time does not belong to any except those which are right? And similarly: Can we discover any characteristic, over and above the mere fact that they are wrong, which belongs to absolutely *all* voluntary actions which are wrong, and which at the same time does not belong to any except those which are wrong? And so, too, in the case of the words "ought" and "duty," it wants to discover some char-

acteristic which belongs to *all* voluntary actions which *ought* to be done or which it is our duty to do, and which does not belong to any except those which we ought to do; and similarly to discover some characteristic which belongs to *all* voluntary actions which ought *not* to be done and which it is our duty *not* to do, and which does not belong to any except these. To all these questions our theory thinks that it can find a comparatively simple answer. And it is this answer which forms the first part of the theory. It is, as I say, a *comparatively* simple answer; but nevertheless it cannot be stated accurately except at some length. And I think it is worth while to try to state it accurately.

To begin with, then, this theory points out that all actions may, theoretically at least, be arranged in a scale, according to the proportion between the *total* quantities of pleasure or pain which they *cause*. And when it talks of the *total* quantities of pleasure or pain which an action causes, it is extremely important to realise that it means quite strictly what it says. We all of us know

that many of our actions do cause pleasure and pain not only to ourselves, but also to other human beings, and sometimes, perhaps, to animals as well ; and that the effects of our actions, in this respect, are often not confined to those which are comparatively direct and immediate, but that their indirect and remote effects are sometimes quite equally important or even more so. But in order to arrive at the *total* quantities of pleasure or pain caused by an action, we should, of course, have to take into account absolutely *all* its effects, both near and remote, direct and indirect ; and we should have to take into account absolutely *all* the beings, capable of feeling pleasure or pain, who were at any time affected by it ; not only ourselves, therefore, and our fellow-men, but also any of the lower animals, to which the action might cause pleasure or pain, however indirectly ; and also any other beings in the Universe, if there should be any, who might be affected in the same way. Some people, for instance, hold that there is a God and that there are disembodied spirits, who may be pleased or pained by our actions ; and, if

this is so, then, in order to arrive at the *total* quantities of pleasure or pain which an action causes, we should have, of course, to take into account, not only the pleasures or pains which it may cause to men and animals upon this earth, but also those which it may cause to God or to disembodied spirits. By the *total* quantities of pleasure or pain which an action causes, this theory means, then, quite strictly what it says. It means the quantities which would be arrived at, if we could take into account absolutely *all* the amounts of pleasure or pain, which result from the action ; no matter how indirect or remote these results may be, and no matter what may be the nature of the beings who feel them.

But if we understand the total quantities of pleasure or pain caused by an action in this strict sense, then obviously, theoretically at least, six different cases are possible. It is obviously theoretically possible in the first place (1) that an action should, in its total effects, cause some pleasure but absolutely no pain ; and it is obviously also possible (2) that, while it causes both pleasure and

pain, the total quantity of pleasure should be *greater* than the total quantity of pain. These are two out of the six theoretically possible cases; and these two may be grouped together by saying that, in both of them, the action in question causes an *excess* of pleasure over pain, or *more* pleasure than pain. This description will, of course, if taken quite strictly, apply only to the second of the two; since an action which causes no pain whatever cannot strictly be said to cause more pleasure than pain. But it is convenient to have some description, which may be understood to cover both cases; and if we describe no pain at all as a *zero* quantity of pain, then obviously we may say that an action which causes some pleasure and no pain, does cause a *greater* quantity of pleasure than of pain, since any positive quantity is greater than zero. I propose, therefore, for the sake of convenience, to speak of both these first two cases as cases in which an action causes an *excess* of pleasure over pain.

But obviously two other cases, which are also theoretically possible, are (1) that in

which an action, in its total effects, causes some pain but absolutely no pleasure, and (2) that in which, while it causes both pleasure and pain, the total quantity of *pain* is greater than the total quantity of *pleasure*. And of both these two cases I propose to speak, for the reason just explained, as cases in which an action causes an *excess* of *pain* over *pleasure*.

There remain two other cases, and two only, which are still theoretically possible; namely (1) that an action should cause absolutely no pleasure and also absolutely no pain, and (2) that, while it causes both pleasure and pain, the total quantities of each should be exactly equal. And in both these two cases, we may, of course, say that the action in question causes *no* excess either of pleasure over pain or of pain over pleasure.

Of absolutely every action, therefore, it must be true, in the sense explained, that it either causes an excess of pleasure over pain, or an excess of pain over pleasure, or neither. This threefold division covers all the six possible cases. But, of course, of any two actions, both of which cause an excess of

pleasure over pain, or of pain over pleasure, it may be true that the excess caused by the one is *greater* than that caused by the other. And, this being so, all actions may, theoretically at least, be arranged in a scale, starting at the top with those which cause the *greatest* excess of pleasure over pain ; passing downwards by degrees through cases where the excess of pleasure over pain is continually smaller and smaller, until we reach those actions which cause no excess either of pleasure over pain or of pain over pleasure : then starting again with those which cause an excess of pain over pleasure, but only the smallest possible one ; going on by degrees to cases in which the excess of pain over pleasure is continually larger and larger ; until we reach, at the bottom, those cases in which the excess of pain over pleasure is the greatest.

The principle upon which this scale is arranged is, I think, perfectly easy to understand, though it cannot be stated accurately except in rather a complicated way. The principle is : That any action which causes an excess of pleasure over pain will always come higher in the scale *either*

than an action which causes a *smaller* excess of pleasure over pain, *or* than an action which causes no excess either of pleasure over pain or of pain over pleasure, *or* than one which causes an excess of pain over pleasure ; That any action, which causes no excess either of pleasure over pain or of pain over pleasure will always come higher than any which causes an excess of pain over pleasure ; and finally That any, which causes an excess of pain over pleasure, will always come higher than one which causes a *greater* excess of pain over pleasure. And obviously this statement is rather complicated. But yet, so far as I can see, there is no simpler way of stating quite accurately the principle upon which the scale is arranged. By saying that one action comes higher in the scale than another, we may mean any one of these five different things ; and I can find no simple expression which will really apply quite accurately to all five cases.

But it has, I think, been customary, among ethical writers, to speak loosely of any action, which comes higher in this scale than another, for any one of these

five reasons, as causing *more* pleasure than that other, or causing a *greater balance* of pleasure over pain. For instance, if we are comparing five different actions, one of which comes higher in the scale than any of the rest, it has been customary to say that, among the five, this is the one which causes a *maximum* of pleasure, or a *maximum balance* of pleasure over pain. To speak in this way is obviously extremely inaccurate, for many different reasons. It is obvious, for instance, that an action which comes lower in the scale may actually produce much more pleasure than one which comes higher, provided this effect is counteracted by its *also* causing a much greater quantity of pain. And it is obvious also that, of two actions, one of which comes higher in the scale than another, *neither* may cause a balance of pleasure over pain, but both actually more pain than pleasure. For these and other reasons it is quite inaccurate to speak as if the place of an action in the scale were determined either by the total quantity of pleasure that it causes, or by the total balance of pleasure over pain. But this way

of speaking, though inaccurate, is also extremely convenient; and of the two alternative expressions, the one which is the most inaccurate is also the most convenient. It is much more convenient to be able to refer to any action which comes higher in the scale as simply causing *more pleasure*, than to have to say, every time, that it causes *a greater balance of pleasure over pain*.

I propose, therefore, in spite of its inaccuracy, to adopt this loose way of speaking. And I do not think the adoption of it need lead to any confusion, provided it is clearly understood, to begin with, that I am going to use the words in this loose way. It must, therefore, be clearly understood that, when, in what follows, I speak of one action as causing more pleasure than another, I shall not mean strictly what I say, but only that the former action is related to the latter in one or other of the five following ways. I shall mean that the two actions are related to one another either (1) by the fact that, while both cause an excess of pleasure over pain, the former causes a greater excess than the latter; or (2) by the fact that, while the former causes

an excess of pleasure over pain, the latter causes no excess whatever either of pleasure over pain, or of pain over pleasure ; or (3) by the fact that, while the former causes an excess of pleasure over pain, the latter causes an excess of pain over pleasure ; or (4) by the fact that, while the former causes no excess whatever either of pleasure over pain or of pain over pleasure, the latter does cause an excess of pain over pleasure ; or (5) by the fact that, while both cause an excess of pain over pleasure, the former causes a smaller excess than the latter. It must be remembered, too, that in every case we shall be speaking of the *total* quantities of pleasure and pain caused by the actions, in the strictest possible sense ; taking into account, that is to say, absolutely *all* their effects, however remote and indirect.

But now, if we understand the statement that one action causes more pleasure than another in the sense just explained, we may express as follows the first principle, which the theory I wish to state lays down with regard to right and wrong, as applied to voluntary actions. This first principle is

a very simple one; for it merely asserts: That a voluntary action is right, whenever and only when the agent could *not*, even if he had chosen, have done any other action instead, which would have caused more pleasure than the one he did do; and that a voluntary action is wrong, whenever and only when the agent *could*, if he had chosen, have done some other action instead, which would have caused more pleasure than the one he did do. It must be remembered that our theory does not assert that any agent ever could have *chosen* any other action than the one he actually performed. It only asserts, that, in the case of all voluntary actions, he *could* have acted differently, *if* he had chosen: not that he could have made the choice. It does not assert, therefore, that right and wrong depend upon what he could *choose*. As to this, it makes no assertion at all: it neither affirms nor denies that they do so depend. It only asserts that they do depend upon what he could have done or could do, *if* he chose. In every case of voluntary action, a man could, *if* he had so chosen just before, have

done at least one other action instead. That was the definition of a voluntary action : and it seems quite certain that many actions are voluntary in this sense. And what our theory asserts is that, where among the actions which he could thus have done instead, *if* he had chosen, there is any one which would have caused more pleasure than the one he did do, then his action is always wrong ; but that in all other cases it is right. This is what our theory asserts, if we remember that the phrase "causing more pleasure" is to be understood in the inaccurate sense explained above.

But it will be convenient, in what follows, to introduce yet another inaccuracy in our statement of it. It asserts, we have seen, that the question whether a voluntary action is right or wrong, depends upon the question whether, among all the other actions, which the agent could have done instead, *if* he had chosen, there is or is not any which would have produced more pleasure than the one he did do. But it would be highly inconvenient, every time we have to mention the theory, to use the whole phrase "all the

other actions which the agent could have done instead, *if* he had chosen." I propose, therefore, instead to call these simply "all the other actions which he *could* have done," or "which were possible to him." This is, of course, inaccurate, since it is, in a sense, not true that he *could* have done them, if he could not have chosen them: and our theory does not pretend to say whether he *ever* could have chosen them. Moreover, even if it is true that he could *sometimes* have chosen an action which he did not choose, it is pretty certain that it is not always so; it is pretty certain that it is *sometimes* out of his power to choose an action, which he certainly could have done, *if* he had chosen. It is not true, therefore, that *all* the actions which he could have done, *if* he had chosen, are actions which in every sense, he *could* have done, even if it is true that some of them are. But nevertheless I propose, for the sake of brevity, to speak of them all as actions which he *could* have done, and this again, I think, need lead to no confusion, if it be clearly understood that I am doing so. It must, then,

e ought. And it would be rash to that such cases never do practically. We all commonly hold that they do : very often indeed we are under no obligation to do one action rather than another ; that it does not matter which we do. We must, then, be careful not to affirm that because it is always our duty to act in a certain way, therefore any particular action, which is always also one which it is our duty to do, is therefore right. This is not so, because, even where an action is right, it does not follow that it is wrong to do something else instead ; if an action is a duty or an action which one positively ought to do, it always follows that it is wrong to do anything else instead. The first consequence, then, which follows, is the distinction between what is right, on the one hand, and what ought to be done or is one's duty, on the other, is that a volun- tary action may be right, without being an action which we ought to do or which it is our duty to do. And from this it follows that the relation between " right " and " ought to be done " is not on a par with that between " wrong " and what

ought *not* to be done. Every action which is wrong is also an action which ought to be done and which it is our duty to do, and also, conversely, every action which ought not to be done, or which it is our duty not to do, is wrong. These three terms are precisely and absolutely co-extensive. To say that an action is wrong, is to imply that it ought not to have been, done; and the converse implication also holds. But in the case of "right" and "ought," only one of the two converse propositions holds. Every action which ought to be done or which it is our duty to do is certainly also right; to say the one is to imply the other. But the converse is not true; since, as we have seen, to say that an action is right is *not* to imply that it ought to be done or that it is our duty to do. An action may be right, without either of these two other things being true of it. To inspect the relation between the positive conceptions "right" and "ought to be done" is not on a par with that between the negative conceptions "wrong" and "ought not to be done." The two positive conceptions

extensive, whereas the two negative ones are so.

And thirdly and finally, it also follows that whereas every voluntary action, without exception, must be either right or wrong, it is by no means necessarily true of every voluntary action that it either ought to be done or ought not to be done,—that it either is our duty to do it, or our duty not to do it.

On the contrary, cases may occur quite frequently where it is neither our duty to do a particular action, nor yet our duty not to do it. This will occur, whenever, the alternatives open to us, there are more, any one of which would be right. And hence we must not suppose that, wherever we have a choice of alternatives before us, there is always some one of them (if we could only find out which is *the* one which we ought to do) while all the rest are definitely wrong. It may quite well be the case that there is more than one among them, which we are under a definite obligation to do, although there always will be at least one which it would be right to do. There will be one which we definitely

only which among actual past voluntary actions *were* right, but also which among those which were possible *would have been* right if they had been done; and not only which among the voluntary actions which actually will be done in the future, *will* be right, but also which among those which will be possible, *would* be right, if they *were* to be done. And in doing this, it does, of course, give us a criterion, or test, or standard, by means of which we could, theoretically at least, discover with regard to absolutely every voluntary action, whichever either has been or will be either actual or possible, whether it was or will be right or wrong. To discover with regard to any voluntary action which was actually done, or which was possible in the past, whether it was right or would have been right, we have only to ask: Could the agent, on the occasion in question, have done anything else which would have produced more utility? If he could, then the action in question was or would have been wrong; if he could not, then it was or would have been right. And similarly, if we want to

the *sole* effect of another were to be the enjoyment of a much more refined one, then they must hold that there would be no reason whatever for preferring the latter to the former, provided only that the mere quantity of pleasure enjoyed in each case were the same. And if the bestial pleasure were ever so slightly more pleasant than the other, then they must say it would be our positive duty to do the action which would bring it about rather than the other. This is a conclusion which does follow from the assertion that actions are right *because* they produce a maximum of pleasure, and which do not follow from the mere assertion that producing a maximum of pleasure is

if we take into account all their further effects, tend to produce more pleasure on the whole than lower ones. There is a good deal to be said for the view that this does actually happen, as the Universe is actually constituted; and that hence an action which causes a higher pleasure to be enjoyed instead of a lower one, will in general cause *more* pleasure in its *total* effects, though it may cause *less* in its *immediate* effects. And this is why those who hold that higher pleasures are in general to be preferred to lower ones, may nevertheless admit that mere quantity of pleasure is always *in itself* a correct *sign* or *criterion* of the rightness of an action.

But those who hold that actions are only right, *because* of the quantity of pleasure they produce, must hold also that, if higher pleasures did not, in their total effects, produce *more* pleasure than lower ones, there would be no reason whatever for preferring them, provided they were not themselves more pleasant. If the *sole* effect of one action were to be the enjoyment of a certain amount of the most bestial or idiotic pleasure, there, and

the *sole* effect of another were to be the enjoyment of a much more refined one, then they must hold that there would be no reason whatever for preferring the latter to the former, provided only that the mere quantity of pleasure enjoyed in each case were the same. And if the bestial pleasure were ever so slightly more pleasant than the other, then they must say it would be our positive duty to do the action which would bring it about rather than the other. This is a conclusion which does follow from the assertion that actions are right *because* they produce a maximum of pleasure, and which does not follow from the mere assertion that the producing a maximum of pleasure is always, *in fact*, a sign of rightness. And it is for this, and similar reasons, that it is im-

the producing of this result does in fact happen to coincide with the producing of other results. They would say that though perhaps, in fact, actual cases never occur in which it *is* or would be wrong to do an action, which produces a maximum of pleasure, it is easy to *imagine* cases in which it *would* be wrong. *If*, for instance, we had to choose between creating a Universe, in which all the inhabitants were capable only of the lowest sensual pleasures, and another in which they were capable of the highest intellectual and æsthetic ones, it would, they would say, plainly be our duty to create the latter rather than the former, even though the mere quantity of pleasure enjoyed in it were rather less than in the former, and still more so if the quantities were equal. Or, to put it shortly, they would say that a world of men is preferable to a world of pigs, even though the pigs might enjoy as much or more pleasure than a world of men. And this is what our theory goes on to deny, when it says that voluntary actions are right, *because* they produce a maximum of pleasure. It implies, by saying this, that actions which

produced a maximum of pleasure *would* always be right, no matter what their effects, in other respects, might be. And hence that it *would* be right to create a world in which there was no intelligence and none of the higher emotions, rather than one in which these were present in the highest degree, provided only that the mere quantity of pleasure enjoyed in the former were ever so little greater than that enjoyed in the latter.

Our theory asserts, then, in its second part, that voluntary actions are right when they are right, *because* they produce a maximum of pleasure; and in asserting this it takes a great step beyond what it asserted in its first part, since it now implies that an action which produced a maximum of pleasure always *would* be right, no matter how its results, in other respects, might compare with those of the other possible alternatives.

But it might be held that, even so, it does not imply that this would be so *absolutely unconditionally*. It might be held that though, in the Universe as actually constituted, actions are right *because* they produce a maximum of pleasure, and hence their right-

ness does not at all depend upon their *other* effects, yet this is only so for some such reason as that, in this Universe, all conscious beings do actually happen to desire pleasure ; but that, if we could imagine a Universe, in which pleasure was not desired, then, in such a Universe, actions would *not* be right because they produced a maximum of pleasure ; and hence that we cannot lay it down absolutely unconditionally that in all conceivable Universes any voluntary action would be right whenever and only when it produced a maximum of pleasure. For some such reason as this, it might be held that ^{we} must distinguish between the mere assertion that voluntary actions are right, when they are right, *because* they produce a maximum of pleasure, and the further assertion that this *would* be so in all conceivable circumstances and in any conceivable Universe. Those who assert the former are by no means necessarily bound to assert the latter also. To assert the latter is to take a still further step.

But the theory I wish to state does, in fact, take this further step. It asserts not only that, in the Universe as it is, voluntary

actions are right *because* they produce a maximum of pleasure, but also that this would be so, *under any conceivable circumstances*: that if any conceivable being, in any conceivable Universe, were faced with a choice between an action which would cause more pleasure and one which would cause less, it would *always* be his duty to choose the former rather than the latter, no matter what the respects might be in which his Universe differed from ours. It may, at first sight, seem unduly bold to assert that any ethical truth can be absolutely unconditional, ⁱⁿ this sense. But many philosophers have held that some fundamental ethical principles certainly are thus unconditional. And a little reflection will suffice to show that the view that they may be so is at all events not absurd. We have many instances of other truths, which seem quite plainly to be of this nature. It seems quite clear, for instance, that it is not only true that twice two do make four, in the Universe as it actually is, but that they necessarily would make four, in any conceivable Universe, no matter how much it might differ from this

one in other respects. And our theory is only asserting that the connection which it believes to hold between rightness and the production of a maximum of pleasure is, in this respect, similar to the connection asserted to hold between the number two and the number four, when we say that twice two are four. It asserts that, if any being whatever, in any circumstances whatever, had to choose between two actions, one of which would produce more pleasure than the other, it always would be his duty to choose the former rather than the latter. That this is absolutely unconditionally true. This assertion obviously goes very ^{far} further, both than the assertion ^{which} it made in its first part, to the effect ^{that} the producing a maximum of pleasure is a *sign* of rightness in the case of all voluntary actions, that ever have been or will be actual or possible, and also than the assertion, that in the Universe, as it is actually constituted, actions are right, when they are right, *because* they produce a maximum of pleasure. But bold as the assertion may seem, it is, at all events, not impossible that we should know it to be true.

Our theory asserts, therefore, in its second part: That, if we had to choose between two actions, one of which would have as its sole or total effects, an effect or set of effects, which we may call A, while the other would have as its sole or total effects, an effect or set of effects, which we may call B, then, if A contained more pleasure than B, it always would be our duty to choose the action which caused A rather than that which caused B. This, it asserts, would be absolutely *always* true, *no matter what A and B might be like in other respects*. And to assert this is, (it ^{is} ~~is~~ ^{which} goes on to say) *equivalent* to asserting ^{set} any effect or set of effects which contains more pleasure is always *intrinsically better* than one which contains less.

By calling one effect or set of effects *intrinsically better* than another it means that it is better *in itself*, quite apart from any accompaniments or further effects which it may have. That is to say: To assert of any one thing, A, that it is *intrinsically better* than another, B, is to assert that if A existed *quite alone*, without any accompaniments

or effects whatever—if, in short, A constituted the whole Universe, it would be better that such a Universe should exist, than that a Universe which consisted solely of B should exist instead. In order to discover whether any one thing is *intrinsically* better than another, we have always thus to consider whether it would be better that the one should exist *quite alone* than that the other should exist *quite alone*. No one thing or set of things, A, ever can be *intrinsically* better than another, B, unless it would be better that A should exist quite alone than that B should exist quite alone. Our theory asserts, therefore, that, wherever it is true that it would be our *duty* to choose A rather than B, if A and B were to be the sole effects of a pair of actions between which we had to choose, there it is always also true that it would be *better* that A should exist quite alone than that B should exist quite alone. And it asserts also, conversely, that wherever it is true that any one thing or set of things, A, is *intrinsically* better than another, B, there it would always also be our duty to choose an action of which A would

be the sole effect rather than one of which B would be the sole effect, if we had to choose between them. But since, as we have seen, it holds that it never could be our duty to choose one action rather than another, unless the total effects of the one contained more pleasure than that of the other, it follows that, according to it, no effect or set of effects, A, can possibly be intrinsically better than another, B, *unless* it contains more pleasure. It holds, therefore, not only that any one effect or set of effects, which contains more pleasure, is always intrinsically better than one which contains less, but also that no effect or set of effects can be intrinsically better than another *unless* it contains more pleasure.

It is plain, then, that this theory assigns a quite unique position to pleasure and pain in two respects; or possibly only in one, since it is just possible that the two propositions which it makes about them are not merely equivalent, but absolutely identical—that is to say, are merely different ways of expressing exactly the same idea. The two propositions are these. (1) That if any one had

to choose between two actions, one of which would, in its total effects, cause more pleasure than the other, it always would be his duty to choose the former; and that it never could be any one's duty to choose one action rather than another, unless its total effects contained more pleasure.

(2) That any Universe, or part of a Universe, which contains more pleasure, is always intrinsically better than one which contains less; and that nothing can be intrinsically better than anything else, unless it contains more pleasure. It does seem to be just possible that these two propositions are merely two different ways of expressing exactly the same idea. The question whether they are so or not simply depends upon the question whether, when we say, "It would be better that A should exist quite alone than that B should exist quite alone," we are or are not saying exactly the same thing, as when we say, "Supposing we had to choose between an action of which A would be the sole effect, and one of which B would be the sole effect, it would be our duty to choose the former rather than the

latter." And it certainly does seem, at first sight, as if the two propositions were not identical; as if we should not be saying exactly the same thing in asserting the one, as in asserting the other. But, even if they are not identical, our theory asserts that they are certainly *equivalent*: that, whenever the one is true, the other is certainly also true. And, if they are not identical, this assertion of equivalence amounts to the very important proposition that: An action is right, only if no action, which the agent could have done instead, would have had intrinsically better results; while an action is wrong, only if the agent *could* have done some other action instead whose total results would have been intrinsically better. It certainly seems as if this proposition were not a mere tautology. And, if so, then we must admit that our theory assigns a unique position to pleasure and pain in two respects, and not in one only. It asserts, first of all, that they have a unique relation to right and wrong; and secondly, that they have a unique relation to *intrinsic value*.

Our theory asserts, then, that any whole

which contains a greater amount of pleasure, is always intrinsically better than one which contains a smaller amount, no matter what the two may be like in other respects; and that no whole can be intrinsically better than another unless it contains more pleasure. But it must be remembered that throughout this discussion, we have, for the sake of convenience, been using the phrase "contains more pleasure" in an inaccurate sense. I explained that I should say of one whole, A, that it contained more pleasure than another, B, whenever A and B were related to one another in either of the five following ways: namely (1) when A and B both contain an excess of pleasure over pain, but A contains a greater excess than B; (2) when A contains an excess of pleasure over pain, while B contains no excess either of pleasure over pain or of pain over pleasure; (3) when A contains an excess of pleasure over pain, while B contains an excess of pain over pleasure, (4) when A contains no excess either of pleasure over pain or of pain over pleasure, while B does contain an excess of pain over pleasure; and (5)

When both A and B contain an excess of pain over pleasure, but A contains a smaller excess than B. Whenever in stating this theory, I have spoken of one whole, or effect, or set of effects, A, as containing more pleasure than another, B, I have always meant merely that A was related to B *in one or other of these five ways*. And so here, when our theory says that every whole which contains a greater amount of pleasure is always intrinsically better than one which contains less, and that nothing can be intrinsically better than anything else unless it contains more pleasure, this must be understood to mean that any whole, A, which stands to another, B, in *any one* of these five relations, is always intrinsically better than B, and that no one thing can be intrinsically better than another, unless it stands to it in *one or other* of these five relations. And it becomes important to remember this, when we go on to take account of another fact.

It is plain that when we talk of one thing being "better" than another we may mean any one of five different things. We may

mean either (1) that while both are positively good, the first is better; or (2) that while the first is positively good, the second is neither good nor bad, but indifferent; or (3) that while the first is positively good, the second is positively bad; or (4) that while the first is indifferent, the second is positively bad; or (5) that while both are positively bad, the first is less bad than the second. We should, in common life, say that one thing was "better" than another, whenever it stood to that other in any one of these five relations. Or, in other words, we hold that among things which stand to one another in the relation of better and worse, some are positively good, others positively bad, and others neither good nor bad, but indifferent. And our theory holds that this is, in fact, the case, with things which have a place in the scale of *intrinsic* value: some of them are intrinsically good, others intrinsically bad, and others indifferent. And it would say that a whole is intrinsically good, whenever and only when it contains an excess of pleasure over pain; intrinsically bad, whenever and only when it contains

an excess of pain over pleasure ; and intrinsically indifferent, whenever and only when it contains neither.

In addition, therefore, to laying down precise rules as to what things are intrinsically *better* or *worse* than others, our theory also lays down equally precise ones as to what things are intrinsically *good* and *bad* and *indifferent*. By saying that a thing is intrinsically good it means that it would be a good thing that the thing in question should exist, even if it existed *quite alone*, without any further accompaniments or effects whatever. By saying that it is intrinsically bad, it means that it would be a bad thing or an evil that it should exist, even if it existed quite alone, without any further accompaniments or effects whatever. And by saying that it is intrinsically indifferent, it means that, if it existed *quite alone*, its existence would be neither a good nor an evil in any degree whatever. And just as the conceptions "intrinsically better" and "intrinsically worse" are connected in a perfectly precise manner with the conceptions "right" and "wrong," so, it maintains, are these

other conceptions also. To say of anything, A, that it is "intrinsically good," is equivalent to saying that, if we had to choose between an action of which A would be the sole or total effect, and an action, which would have absolutely no effects at all, it would always be our duty to choose the former, and wrong to choose the latter. And similarly to say of anything, A, that it is "intrinsically bad," is equivalent to saying that, if we had to choose between an action of which A would be the sole effect, and an action which would have absolutely no effects at all, it would always be our duty to choose the latter and wrong to choose the former. And finally, to say of anything, A, that it is "intrinsically indifferent," is equivalent to saying that, if we had to choose between an action, of which A would be the sole effect, and an action which would have absolutely no effects at all, it would not matter which we chose: either choice would be equally right.

To sum up, then, we may say that, in its second part, our theory lays down three principles. It asserts (1) that anything

whatever, whether it be a single effect, or a whole set of effects, or a whole Universe, is *intrinsically good*, whenever and only when it either is or contains an excess of pleasure over pain; that anything whatever is *intrinsically bad*, whenever and only when it either is or contains an excess of pain over pleasure; and that all other things, no matter what their nature may be, are intrinsically indifferent. It asserts (2) that any one thing, whether it be a single effect, or a whole set of effects, or a whole Universe, is intrinsically *better* than another, whenever and only when the two are related to one another in ⁴one or other of the five following ways: namely, when either (a) while both are intrinsically good, the second is not so good as the first; or (b) while the first is intrinsically good, the second is intrinsically indifferent; or (c) while the first is intrinsically good, the second is intrinsically bad; or (d) while the first is intrinsically indifferent, the second is intrinsically bad; or (e) while both are intrinsically bad, the first is not so bad as the second. And it asserts (3) that, if we had to choose between two actions

one of which would have intrinsically better total effects than the other, it always would be our duty to choose the former, and wrong to choose the latter; and that no action ever can be right *if* we could have done anything else instead which would have had intrinsically better total effects, nor wrong, *unless* we could have done something else instead which would have had intrinsically better total effects. From these three principles taken together, the whole theory follows. And whether it be true or false, it is, I think, at least a perfectly clear and intelligible theory. Whether it is or is not of any practical importance, is, indeed, another question. But, even if it were of none whatever, it certainly lays down propositions of so fundamental and so far-reaching a character, that it seems worth while to consider whether they are true or false. There remain, I think, only two points which should be noticed with regard to it, before we go on to consider the principal objections which may be urged against it.

It should be noticed, first, that, though

this theory asserts that nothing is *intrinsically* good, unless it is or contains an excess of pleasure over pain, it is very far from asserting that nothing is *good*, unless it fulfils this condition. By saying that a thing is *intrinsically good*, it means, as has been explained, that the existence of the thing in question *would* be a good, even if it existed quite alone, without any accompaniments or effects whatever; and it is quite plain that when we call things "good" we by no means always mean this: we by no means always mean that they *would* be good, even if they existed quite alone. Very often, for instance, when we say that a thing is "good," we mean that it is good *because of its effects*; and we should not for a moment maintain that it *would* be good, even if it had no effects at all. We are, for instance, familiar with the idea that it is sometimes a good thing for people to suffer pain; and yet we should be very loth to maintain that in all such cases their suffering *would* be a good thing, even if nothing were gained by it—if it had no further effects. We do, in general, maintain that suffering is good,

only *where* and *because* it has further good effects. And similarly with many other things. Many things, therefore, which are *not* "intrinsically" good, may nevertheless be "good" in some one or other of the senses in which we use that highly ambiguous word. And hence our theory can and would quite consistently maintain that, while nothing is *intrinsically* good except pleasure or wholes which contain pleasure, many other things really are "good"; and similarly that, while nothing is *intrinsically* bad except pain or wholes which contain it, yet many other things are really "bad." It would, for instance, maintain that it is *always* a good thing to act rightly, and a bad thing to act wrongly; although it would say at the same time that, since actions, strictly speaking, do not *contain* either pleasure or pain, but are only accompanied by or causes of them, a right action is *never intrinsically* good, nor a wrong one *intrinsically* bad. And similarly it would maintain that it is perfectly true that some men are "good," and others "bad," and some better than others; although no man can strictly be

said to *contain* either pleasure or pain, and hence none can be either intrinsically good or intrinsically bad or intrinsically better than any other. It would even maintain (and this also it can do quite consistently), that events which are *intrinsically* good are nevertheless very often bad, and intrinsically bad ones good. It would, for instance, say that it is often a very bad thing for a man to enjoy a particular pleasure on a particular occasion, although the event, which consists in his enjoying it, may be intrinsically good, since it contains an excess of pleasure over pain. It may often be a very bad thing that such an event should happen, because it *causes* the man himself or other beings to have less pleasure or more pain in the future, than they would otherwise have had. And for similar reasons it may often be a very good thing that an intrinsically bad event should happen.

It is important to remember all this, because otherwise the theory may appear much more paradoxical than it really is. It may, for instance, appear, at first sight, as if it denied all value to anything except pleasure and wholes which contain it—a view which

would be extremely paradoxical if it were held. But it does *not* do this. It does not deny all value to other things, but only all *intrinsic* value—a very different thing. It only says that none of them *would* have any value if they existed quite alone. But, of course, as a matter of fact, none of them do exist quite alone, and hence it may quite consistently allow that, as it is, many of them do have very great value. Concerning kinds of value, other than intrinsic value, it does not profess to lay down any general rules at all. And its reason for confining itself to intrinsic value is because it holds that this and this alone is related to right and wrong in the perfectly definite manner explained above. Whenever an action is right, it is right only if and because the total effects of no action, which the agent could have done instead, would have had more *intrinsic* value; and whenever an action is wrong, it is wrong only if and because the total effects of some other action, which the agent could have done instead, would have had more *intrinsic* value. This proposition, which is true of

intrinsic value, is not, it holds, true of value of any other kind.

And a second point which should be noticed about this theory is the following. It is often represented as asserting that pleasure is the only thing which is *ultimately* good or desirable, and pain the only thing which is *ultimately* bad or undesirable; or as asserting that pleasure is the only thing which is good *for its own sake*, and pain the only thing which is bad *for its own sake*. And there is, I think, a sense in which it does assert this. But these expressions are not commonly carefully defined; and it is worth noticing that, if our theory does assert these propositions, the expressions "*ultimately* good" or "good *for its own sake*" must be understood in a different sense from that which has been assigned above to the expression "*intrinsically* good." We must not take "*ultimately* good" or "good *for its own sake*" to be synonyms for "*intrinsically* good." For our theory most emphatically does *not* assert that pleasure is the only thing *intrinsically* good, and pain the only thing *intrinsically* evil. On the contrary,

it asserts that any whole which *contains* an excess of pleasure over pain is *intrinsically* good, no matter how much else it may contain besides; and similarly that any whole which contains an excess of pain over pleasure is *intrinsically* bad. This distinction between the conception expressed by "ultimately good" or "good for its own sake," on the one hand, and that expressed by "*intrinsically* good," on the other, is not commonly made; and yet obviously we must make it, if we are to say that our theory does assert that pleasure is the only *ultimate* good, and pain the only *ultimate* evil. The two conceptions, if used in this way, have one important point in common, namely, that both of them will only apply to things whose existence *would* be good, even if they existed quite alone. Whether we assert that a thing is "ultimately good" or "good for its own sake" or "*intrinsically* good," we are always asserting that it would be good, even if it existed quite alone. But the two conceptions differ in respect of the fact that, whereas a whole which is "*intrinsically* good" may contain parts which

are *not* intrinsically good, *i.e.* *would* not be good, if they existed quite alone; anything which is "ultimately good" or "good for its own sake" can contain no such parts. This, I think, is the meaning which we must assign to the expressions "ultimately good" or "good for its own sake," if we are to say that our theory asserts pleasure to be the *only* thing "ultimately good" or "good for its own sake." We may, in short, divide intrinsically good things into two classes: namely (1) those which, while as wholes they are intrinsically good, nevertheless contain some parts which are not intrinsically good; and (2) those, which either have no parts at all, or, if they have any, have none but what are themselves intrinsically good. And we may thus, if we please, confine the terms "ultimately good" or "good for their own sakes" to things which belong to the second of these two classes. We may, of course, make a precisely similar distinction between two classes of intrinsically bad things. And it is only if we do this that our theory can be truly said to assert that nothing is "ultimately

good" or "good for its own sake," except pleasure; and nothing "ultimately bad" or "bad for its own sake," except pain.

Such is the ethical theory which I have chosen to state, because it seems to me particularly simple, and hence to bring out particularly clearly some of the main questions which have formed the subject of ethical discussion.

What is specially important is to distinguish the question, which it professes to answer in its first part, from the much more radical questions, which it professes to answer in its second. In its first part, it ~~only~~ professes to answer the question: What characteristic is there which does actually, *as a matter of fact*, belong to all right voluntary actions, which ever have been or will be done in this world? While, in its second part, it professes to answer the much more fundamental question: What characteristic is there which *would* belong to absolutely any voluntary action, which was right, in any conceivable Universe, and under any conceivable circumstances? These two questions are obviously extremely different, and by the theory I have

stated I mean a theory which does profess to give an answer to *both*.

Whether this theory has ever been held in exactly the form in which I have stated it, I should not like to say. But many people have certainly held something very like it; and it seems to be what is *often* meant by the familiar name "Utilitarianism," which is the reason why I have chosen this name as the title of these two chapters. It must not, however, be assumed that anybody who talks about "Utilitarianism" *always* means precisely this theory in all its details. On the contrary, many even of those who call themselves Utilitarians would object to some of its most fundamental propositions. One of the difficulties which occurs in ethical discussions is that no single name, which has ever been proposed as the name of an ethical theory, has any absolutely fixed significance. On the contrary, every name may be, and often is, used as a name for several different theories, which may differ from one another in very important respects. Hence, whenever anybody uses such a name, you can never trust to the

name alone, but must always look carefully to see exactly what he means by it. For this reason I do not propose, in what follows, to give any name at all to this theory which I have stated, but will refer to it simply as the theory stated in these first two chapters.

CHAPTER III

THE OBJECTIVITY OF MORAL JUDGMENTS

AGAINST the theory, which has been stated in the last two chapters, an enormous variety of different objections may be urged; and I cannot hope to deal with nearly all of them. What I want to do is to choose out those, which seem to me to be the most important, because they are the most apt to be strongly felt, and because they concern extremely general questions of principle. It seems to me that some of these objections are well founded and that others are not, according as they are directed against different parts of what our theory asserts. And I propose, therefore, to split up the theory into parts, and to consider separately the chief objections which might be urged against each of these different parts.

And we may begin with an extremely