

EDITOR OF
*The Home University Library
of Modern Knowledge*
SIR GEORGE CLARK, D.LITT., F.B.A.

ETHICS

G. E. MOORE
O.M., LL.D., F.B.A.

LONDON
OXFORD UNIVERSITY PRESS
NEW YORK TORONTO

Oxford University Press, Amen House, London E.C.4

GLASGOW NEW YORK TORONTO MELBOURNE WELLINGTON

BOMBAY CALCUTTA MADRAS KARACHI KUALA LUMPUR

CAPE TOWN IBADAN NAIROBI ACCRA

First edition 1912

Reprinted 1925 (twice), 1927, 1928, 1930,

1936, 1939 and 1944

Reset 1947

Reprinted 1949, 1952, 1955, 1958 and 1961

CONTENTS

I	UTILITARIANISM	7
II	UTILITARIANISM (<i>concluded</i>)	27
III	THE OBJECTIVITY OF MORAL JUDGEMENTS	50
IV	THE OBJECTIVITY OF MORAL JUDGEMENTS (<i>concluded</i>)	83
V	RESULTS THE TEST OF RIGHT AND WRONG	106
VI	FREE WILL	122
VII	INTRINSIC VALUE	138
	NOTE ON BOOKS	156
	INDEX	159

PRINTED IN GREAT BRITAIN

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 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 even if we cannot be justified either in 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 realize 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 characteristic 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

realize 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, be clearly understood that, when, in what follows, I speak of all the actions which the agent could have done, or all those open to him under the circumstances, I shall mean only all those which he could have done, *if* he had chosen.

Understanding this, then, we may state the first principle which our theory lays down quite briefly by saying: 'A voluntary action is right, whenever and only when no other action possible to the agent under the circumstances would have caused more pleasure; in all other cases, it is wrong.' This is its answer to the questions: What characteristic is there which belongs to *all* voluntary actions which are right, and *only* those among them which are right? and what characteristic is there which belongs to *all* those which are wrong, and *only* to those which are wrong? But it also asked the very same questions with regard to two other classes of voluntary actions—those which *ought* or *ought not* to be done, and those which it is our *duty* to do or not to do. And its answer to the question concerning these conceptions differs from its answer to the question concerning right and wrong in a way, which is, indeed, comparatively unimportant, but which yet deserves to be noticed.

It may have been observed that our theory does *not* assert that a voluntary action is right only where it causes *more* pleasure than any action which the agent could have done instead. It confines itself to asserting that, in order to be right, such an action must cause at least *as much* pleasure as any which the agent could have done instead. And it confines itself in this way for the following reason. It is obviously possible, theoretically at least, that, among the alternatives open to an agent at a given moment, there may be two or more which would produce precisely *equal* amounts of pleasure, while all of them produced more than any of the other possible alternatives; and in such cases, our theory would say, *any one* of these actions would be perfectly right. It recognizes, therefore, that there may be cases in which no single one of the actions open to the agent can be distinguished as *the* right one to do; that in many cases, on the contrary, several different actions may all be equally right; or, in other words, that to say that a man acted rightly does not necessarily imply that, if he had done anything else instead, he would have acted wrongly. And this is certainly in accordance with common usage. We all do constantly imply that sometimes when a man was right in doing what he did, yet he might have been equally right, if he had acted differently: that there may be several different alternatives open to him, none of which can definitely be said to be wrong. This is why our theory refuses to commit itself to the view that an action is right only where it produces *more* pleasure than any of the other possible alternatives. For, if this were so; then it would follow that no two alternatives could ever be *equally* right: some one of them would always have to be *the* right one, and all the rest wrong. But it is precisely in this respect that it holds that the concep-

tions of 'ought' and of 'duty' differ from the conception of what is 'right'. When we say that a man 'ought' to do one particular action, or that it is his 'duty' to do it, we do imply that it would be wrong for him to do *anything* else.) And hence our theory holds that, in the case of 'ought' and 'duty' we may say, what we could not say in the case of 'right', namely, that an action ought to be done or is our duty, only where it produces *more* pleasure than any which we could have done instead.

From this distinction several consequences follow. It follows firstly that a voluntary action may be 'right' without being an action which we 'ought' to do or which it is our 'duty' to do. It is, of course, always our duty to act rightly, in the sense that, if we don't act rightly, we shall always be doing what we ought not. It is, therefore, true, in a sense, that whenever we act rightly, we are always doing our duty and doing what we ought. But what is not true is that, whenever a particular action is right, it is always our duty to do that particular action and no other. This is not true, because, theoretically at least, cases may occur in which some other action would be quite equally right, and in such cases, we are obviously under no obligation whatever to do the one rather than the other: whichever we do, we shall be doing our duty and doing as we ought. And it would be rash to affirm that such cases never do practically occur. We all commonly hold that they do: that very often indeed we are under no positive obligation to do one action rather than some other; 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 rightly, therefore any particular action, which is right, is always also one which it is our duty to do. This is not so, because, even where an

action is right, it does not follow that it would be wrong to do something else instead; whereas, if an action is a duty or an action which we positively ought to do, it always would be wrong to do anything else instead.

The first consequence, then, which follows, from this distinction between what is right, on the one hand, and what ought to be done or is our duty, on the other, is that a voluntary 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 further that the relation between 'right' and what 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 not to be done and which it is our duty not 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 negative terms are precisely and absolutely coextensive. To say that an action is or was wrong, is to imply that it ought not to be, or 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 is our duty, is certainly also right; to say the one thing of any action is to imply the other. But here 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: an action may be right, without either of these two other things being true of it. In this respect 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 are not coextensive, 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, among the alternatives open to us, there are two or more, any one of which would be equally right. And hence we must not suppose that, wherever we have a choice of actions before us, there is always some one among them (*if* we could only find out which), 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 no one among them, which we are under a positive obligation to do, although there always must be at least one which it would be right to do. There will be one which we definitely *ought* to do, in those cases and those cases only, where there happens to be *only* one which is right under the circumstances—where, that is to say, there are not several which would all be equally right, but some one of the alternatives open to us is *the* only right thing to do. And hence in many cases we cannot definitely say of a voluntary action either that it was the agent's duty to do it nor yet that it was his duty not to do it. There may be cases in which none of the alternatives open to us is definitely prescribed by duty.

To sum up, then: The answers which this theory gives to its first set of questions is as follows. A characteristic which belongs to all right voluntary actions, and only to those which are right, is, it says, this: That they all cause at least *as much* pleasure as any action which the agent could have done instead; or, in other

words, they all produce *a* maximum of pleasure. A characteristic which belongs to all voluntary actions, which *ought* to be done or which it is our *duty* to do, and only to these, is, it says, the slightly different one: That they all cause *more* pleasure than any which the agent could have done instead; or, in other words, among all the possible alternatives, it is they which produce *the* maximum of pleasure. And finally, a characteristic which belongs to all voluntary actions which are wrong, or which ought not to be done, or which it is our duty not to do, and which belongs only to these, is, in all three cases the same, namely: That they all cause *less* pleasure than some other action which the agent could have done instead. These three statements together constitute what I will call the first part of the theory; and, whether we agree with them or not, it must, I think, at least be admitted that they are propositions of a very fundamental nature and of a very wide range, so that it would be worth while to know, if possible, whether they are true.

But this first part of the theory is by no means the whole of it. There are two other parts of it, which are at least equally important; and, before we go on to consider the objections which may be urged against it, it will, I think, be best to state these other parts. They may, however, conveniently form the subject of a new chapter.

CHAPTER II

UTILITARIANISM (*concluded*)

IN the last chapter I stated the first part of an ethical theory, which I chose out for consideration, not because I agreed with it, but because it seemed to me to bring out particularly clearly the distinction between some of the most fundamental subjects of ethical discussion. This first part consisted in asserting that there is a certain characteristic which belongs to absolutely *all* voluntary actions which are right, and *only* to those which are right; another closely allied characteristic which belongs to *all* voluntary actions which ought to be done or are duties, and *only* to these; a third characteristic which belongs to *all* voluntary actions which are wrong, ought not to be done, or which it is our duty not to do, and *only* to those voluntary actions of which these things are true. And when the theory makes these assertions it means the words 'all' and 'only' to be understood quite strictly. That is to say, it means its propositions to apply to absolutely every voluntary action, which ever has been done, or ever will be done, no matter who did it, or when it was or will be done; and not only to those which actually have been or will be done, but also to all those which have been or will be *possible*, in a certain definite sense.

The sense in which it means its propositions to apply to *possible*, as well as actual, voluntary actions, is, it must be remembered, only if we agree to give the name 'possible' to all those actions which an agent *could* have done, *if* he had chosen, and to those which, in the future, any agent will be able to do, *if* he were