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

to choose to do them. Possible actions, in this sense, form a perfectly definite group; and we do, as a matter of fact, often make judgements as to whether they would have been or would be right, and as to whether they ought to have been done in the past, or ought to be done in the future. We say, 'So-and-so ought to have done this on that occasion', or 'It would have been perfectly right for him to have done this', although as a matter of fact, he did not do it; or we say, 'You ought to do this', or 'It will be quite right for you to do this', although it subsequently turns out, that the action in question is one which you do not actually perform. Our theory says, then, with regard to all actions, which were in this sense possible in the past, that they *would have been* right, if and only if they *would* have produced a maximum of pleasure; just as it says that all actual past voluntary actions *were* right, if and only if they *did* produce a maximum of pleasure. And similarly, with regard to all voluntary actions which will be possible in the future, it says that they will be right, if and only if they *would* produce a maximum of pleasure; just as it says with regard to all that will actually be done, that they will be right, if and only if they *do* produce a maximum of pleasure.

Our theory does, then, even in its first part, deal, in a sense, with possible actions, as well as actual ones. It professes to tell us, not 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 not. If we want to discover with regard to a voluntary action which was actually done or 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 instead, which would have produced more pleasure? 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 discover with regard to an action, which we are contemplating in the future, whether it would be right for us to do it, we have only to ask: Could I do anything else instead which would produce more pleasure? If I could, it will be wrong to do the action; if I could not, it will be right. Our theory does then, even in its first part, profess to give us an absolutely universal *criterion* of right and wrong; and similarly also an absolutely universal *criterion* of what ought or ought not to be done.

But though it does this, there is something else which it does not do. It only asserts, in this first part, that the producing of a maximum of pleasure is a characteristic, which did and will belong, *as a matter of fact*, to all right voluntary actions (actual or possible), and only to right ones; it does not, in its first part, go on to assert that it is *because* they possess this characteristic that such actions are right. This second assertion is the first which it goes on to make in its second part; and everybody can see, I think, that there is an important difference between the two assertions.

Many people might be inclined to admit that, when-

ever a man acts wrongly, his action always does, on the whole, result in greater unhappiness than would have ensued if he had acted differently; and that when he acts rightly this result *never* ensues: that, on the contrary, right action always does in the end bring about at least as much happiness, on the whole, as the agent could possibly have brought about by any other action which was in his power. The proposition that wrong action always *does*, and (considering how the Universe is constituted) always *would*, in the long run, lead to less pleasure than the agent could have brought about by acting differently, and that right action never *does* and never *would* have this effect, is a proposition which a great many people might be inclined to accept; and this is all which, in its first part, our theory asserts. But many of those who would be inclined to assent to this proposition, would feel great hesitation in going on to assert that this is *why* actions are right or wrong respectively. There seems to be a very important difference between the two positions. We may hold, for instance, that an act of murder, whenever it is wrong, always does produce greater unhappiness than would have followed if the agent had chosen instead some one of the other alternatives, which he could have carried out, *if* he had so chosen; and we may hold that this is true of all other wrong actions, actual or possible, and never of any right ones: but it seems a very different thing to hold that murder and all other wrong actions are wrong, when they are wrong, *because* they have this result—*because* they produce less than the possible maximum of pleasure. We may hold, that is to say, that the fact that it does produce or would produce *less* than a maximum of pleasure is absolutely always a *sign* that a voluntary action is wrong, while the fact that it does produce or

would produce a maximum of pleasure is absolutely always a *sign* that it is right; but this does not seem to commit us to the very different proposition that these results, besides being *signs* of right and wrong, are also the *reasons* why actions are right when they are right, and wrong when they are wrong. Everybody can see, I think, that the distinction is important; although I think it is often overlooked in ethical discussions. And it is precisely this distinction which separates what I have called the first part of our theory, from the first of the assertions which it goes on to make in its second part. In its first part it only asserts that the producing or not producing a maximum of pleasure are, absolutely universally, *signs* of right and wrong in voluntary actions; in its second part it goes on to assert that it is *because* they produce these results that voluntary actions are right when they are right, and wrong when they are wrong.

There is, then, plainly some important difference between the assertion, which our theory made in its first part, to the effect that all right voluntary actions, and only those which are right, do, *in fact*, produce a maximum of pleasure, and the assertion, which it now goes on to make, that this is *why* they are right. And if we ask why the difference is important, the answer is, so far as I can see, as follows. Namely, if we say that actions are right, *because* they produce a maximum of pleasure, we imply that, provided they produced this result, they *would* be right, *no matter what other effects they might produce* as well. We imply, in short, that their rightness does *not* depend at all upon their other effects, but *only* on the quantity of pleasure that they produce. And this is a very different thing from merely saying that the producing a maximum of pleasure is always, as a matter of fact, a *sign* of rightness.

It is quite obvious, that, in the Universe as it is actually constituted, pleasure and pain are by no means the only results of any of our actions: they all produce immense numbers of other results as well. And so long as we merely assert that the producing a maximum of pleasure is a *sign* of rightness, we leave open the possibility that it is so only because this result does always, as a matter of fact, happen to coincide with the production of *other* results; but that it is partly upon these other results that the rightness of the action depends. But so soon as we assert that actions are right, *because* they produce a maximum of pleasure, we cut away this possibility; we assert that actions which produced such a maximum *would* be right, even if they did not produce any of the other effects, which, as a matter of fact, they always do produce. And this, I think, is the chief reason why many persons who would be inclined to assent to the first proposition, would hesitate to assent to the second.

It is, for instance, commonly held that some pleasures are higher or better than others, even though they may not be more pleasant; and that where we have a choice between procuring for ourselves or others a higher or a lower pleasure, it is generally right to prefer the former, even though it may perhaps be less pleasant. And, of course, even those who hold that actions are only right because of the quantity of pleasure they produce, and not at all because of the quality of these pleasures, might quite consistently hold that it is *as a matter of fact* generally right to prefer higher pleasures to lower ones, even though they may be less pleasant. They might hold that this is the case, on the ground that higher pleasures, even when less pleasant in themselves, do, 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 fact*, 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, then 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, 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 important to distinguish the two propositions.

To many persons it may seem clear that it *would* be our duty to prefer some pleasures to others, even if they did not entail a greater *quantity* of pleasure; and hence that though actions which produce a maximum of pleasure are perhaps, *in fact*, always right, they are not right *because* of this, but only because 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 aesthetic 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 rightness 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 connexion which it believes to hold between rightness and the production of a maximum of pleasure is, in this respect, similar to the connexion 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 much 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 now goes on to say) *equivalent* to asserting that 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, when-

ever 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 one 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.