CHAPTER VII

INTRINSIC VALUE

THE main conclusions, at which we have arrived so far with regard to the theory stated in Chapters I and II, may be briefly summed up as follows. I tried to show, first of all, (1) that to say that a voluntary action is *right*, or *ought* to be done, or is *wrong*, is *not* the same thing as to say that any being or set of beings whatever, either human or non-human, has towards it any mental attitude whatever—either an attitude of feeling, or of willing, or of thinking something about it; and that hence no proof to the effect that any beings, human or non-human, have any such attitude towards an action is sufficient to show that it is right, or ought to be done, or is wrong; and (2) similarly, that to say that any one thing or state of things is *intrinsically good*, or *intrinsically bad*, or that one is *intrinsically better* than another, is also not the same thing as to say that any being or set of beings has towards it any mental attitude whatever—either an attitude of feeling, or of desiring, or of thinking something about it; and hence that here again no proof to the effect that any being or set of beings *has* some such mental attitude towards a given thing or state of things is ever sufficient to show that it is intrinsically good or bad. These two points are extremely important, because the contrary view is very commonly held, in some form or other, and because (though this is not always seen), whatever form it be held in, it is absolutely fatal to one or both of two very fundamental principles, which our theory implies. In many of their forms such views are fatal to the principle (1) that no action is ever *both* right and wrong;

and hence also to the view that there is any characteristic whatever which *always* belongs to right actions and *never* to wrong ones; and in *all* their forms they are fatal to the principle, (2) that if it is once the duty of any being to do an action whose total effects will be A rather than one whose total effects will be B, it must *always* be the duty of any being to do an action whose total effects will be precisely similar to A rather than one whose total effects will be precisely similar to B, if he has to choose between them.

I tried to show, then, first of all, that these two principles may be successfully defended against this first line of attack—the line of attack which consists in saying (to put it shortly) that 'right' and 'good' are merely *subjective* predicates. But we found next that even those who admit and insist (as many do) that 'right' and 'intrinsically good' are *not* subjective predicates, may yet attack the second principle on another ground. For this second principle implies that the question whether an action is right or wrong must always depend upon its *actual* consequences; and this view is very commonly disputed on one or other of three grounds, namely (1) that it sometimes depends merely on the *intrinsic nature* of the action, or, in other words, that certain kinds of actions would be absolutely always right, and others absolutely always wrong, *whatever* their consequences might be, or (2) that it depends, partly or wholly, on the *motive* from which the action is done, or (3) that it depends on the question whether the agent had *reason to expect* that its consequences would be the best possible. I tried, accordingly, to show next that each of these three views is untrue.

But, finally, we raised, in the last chapter, a question as to the *precise* sense in which right and wrong do

depend upon the actual consequences. And here for the first time we came upon a point as to which it seemed very doubtful whether our theory was right. All that could be agreed upon was that a voluntary action is right whenever and only when its total consequences are *as* good, intrinsically, as any that would have followed from any action which the agent *could have* done instead. But we were unable to arrive at any certain conclusion as to the precise sense in which the phrase '*could have*' must be understood if this proposition is to be true; and whether, therefore, it *is* true, if we give to these words the precise sense which our theory gave to them.

I conclude, then, that the theory stated in Chapters I and II is right so far as it merely asserts the three principles (1) That there *is* some characteristic which belongs and must belong to absolutely *all* right voluntary actions and to *no* wrong ones; (2) That one such characteristic consists in the fact that the total consequences of right actions must always be as good, intrinsically, as any which it was *possible* for the agent to produce under the circumstances (it being uncertain, however, in what sense precisely the word 'possible' is to be understood), whereas this can never be true of wrong ones; and (3) That if any set of consequences A is once intrinsically better than another set B, any set precisely similar to A must always be intrinsically better than a set precisely similar to B. We have, indeed, not considered all the objections which might be urged against these three principles; but we have, I think, considered all those which are most commonly urged, *with one single exception.* And I must now briefly state what this one remaining objection is, before I go on to point out the respect in which this theory, which was stated in Chapters I

and II, seems to me to be utterly wrong, in spite of being right as to all these three points.

This one last objection may be called the objection of Egoism; and it consists in asserting that no agent can ever be under any obligation to do the action, whose *total* consequences will be the best possible, *if* its total effects upon *him*, personally, are not the best possible; or in other words that it always would be *right* for an agent to choose the action whose total effects *upon himself* would be the best, even if *absolutely all* its effects (taking into account its effects on other beings as well) would *not* be the best. It asserts in short that it can never be the duty of any agent to sacrifice his own good to the general good. And most people, who take this view, are, I think, content to assert this, without asserting further that it must always be his positive *duty* to prefer his own good to the general good. That is to say, they will admit that a man may be acting *rightly*, even if he *does* sacrifice his own good to the general good; they only hold that he will be acting *equally* rightly, if he does *not*. But there are some philosophers who seem to hold that it must *always* be an agent's positive duty to do what is best for *himself—always*, for instance, to do what will conduce most to his own 'perfection', or his own salvation, or his own 'self-realization'; who imply, therefore, that it would be his duty so to act, even if the action in question did *not* have the best possible consequences upon the whole.

Now the question, whether this view is true, in either of these two different forms, would, of course, be of no practical importance, if it were true that, as a matter of fact, every action which most promotes the general good always *also* most promotes the agent's own good, and vice versa. And many philosophers have

taken great pains to try to show that this *is* the case: some have even tried to show that it *must* necessarily be the case. But it seems to me that none of the arguments which have been used to prove this proposition really do show that it is by any means *universally* true. A case, for instance, may arise in which, if a man is to secure the best consequences for the world as a whole, it may be absolutely necessary that he should sacrifice his own life. And those who maintain that, even in such a case, he will absolutely always be securing the greatest possible amount of good *for himself*, must either maintain that in some future life he will receive goods sufficient to compensate him for all that he might have had during many years of continued life in this world—a view to which there is the objection that it may be doubted, whether we shall have any future life at all, and that it is even more doubtful, what, *if* we shall, that life will be like; or else they must maintain the following paradox.

Suppose there are two men, A and B, who up to the age of thirty have lived lives of equal intrinsic value; and that at that age it becomes the duty of each of them to sacrifice his life for the general good. Suppose A does his duty and sacrifices his life, but B does not, and continues to live for thirty years more. Those who hold that the agent's own good *always* coincides with the general good, must then hold that B's sixty years of life, no matter how well the remaining thirty years of it may be spent, cannot possibly have so much intrinsic value as A's thirty years. And surely this is an extravagant paradox, however much intrinsic value we may attribute to those final moments of A's life in which he does his duty at the expense of his life; and however high we put the loss in intrinsic value to B's life, which arises from the fact that, in this one

instance, he failed to do his duty. B may, for instance, repent of this one act and the whole of the remainder of his life may be full of the highest goods; and it seems extravagant to maintain that all the goods there may be in this last thirty years of it cannot possibly be enough to make his life more valuable, intrinsically, than that of A.

I think, therefore, we must conclude that a maximum of true good, for ourselves, is by no means always secured by those actions which are necessary to secure a maximum of true good for the world as a whole; and hence that it *is* a question of practical importance, whether, in such cases of conflict, it is always a duty, or right, for us to prefer our own good to the general good. And this is a question which, so far as I can see, it is impossible to decide by argument one way or the other. If any person, after clearly considering the question, comes to the conclusion that he can never be under any obligation to sacrifice his own good to the general good, if they *were* to conflict, or even that it would be wrong for him to do so, it is, I think, impossible to prove that he is mistaken. But it is certainly equally impossible for him to prove that he is not mistaken. And, for my part, it seems to me quite self-evident that he is mistaken. It seems to me quite self-evident that it must always be our duty to do what will produce the best effects *upon the whole*, no matter how bad the effects upon ourselves may be and no matter how much good we ourselves may lose by it.

I think, therefore, we may safely reject this last objection to the principle that it must always be the duty of every agent to do that one, among all the actions which he *can* do on any given occasion, whose *total consequences* will have the greatest intrinsic value;

and we may conclude, therefore, that the theory stated in Chapters I and II is right as to all the three points yet considered, except for the doubt as to the precise sense in which the words 'can do' are to be understood in this proposition. But obviously on any theory which maintains, as this one does, that right and wrong depend on the intrinsic value of the consequences of our actions, it is extremely important to decide rightly what kinds of consequences *are* intrinsically better or worse than others. And it is on this important point that the theory in question seems to me to take an utterly wrong view. It maintains, as we saw in Chapter II, that any whole which contains *more pleasure* is always intrinsically better than one which contains less, and that none can be intrinsically better, *unless* it contains more pleasure; it being remembered that the phrase 'more pleasure', in this statement, is not to be understood as meaning strictly what it says, but as standing for any one of five different alternatives, the nature of which was fully explained in our first two chapters. And the last question we have to raise is, therefore: Is this proposition true or not? and if not, what *is* the right answer to the question: What kinds of things are intrinsically better or worse than others?

And first of all it is important to be quite clear as to how this question is related to another question, which is very liable to be confused with it: namely the question whether the proposition which was distinguished in Chapter I, as forming *the first part* of the theory there stated, is true or not: I mean, the proposition that quantity of pleasure is a correct *criterion* of right and wrong, or that, *in this world*, it always is, *as a matter of fact*, our duty to do the action which will produce a maximum of pleasure, or (for

this is, perhaps, more commonly held) to do the action which, *so far as we can see*, will produce such a maximum. This latter proposition has been far more often *expressly* held than the proposition that what contains more pleasure is *always* intrinsically better than what contains less; and many people may be inclined to think they are free to maintain it, even if they deny that the intrinsic value of every whole is *always* in proportion to the quantity of pleasure it contains. And so, *in a sense*, they are; for it is quite possible, *theoretically*, that quantity of pleasure should always be a correct *criterion* of right and wrong, here in this world, even if intrinsic value is not always in exact proportion to quantity of pleasure. But though this is theoretically possible, it is, I think, easy to see that it is extremely *unlikely* to be the case. For if it were the case, what it would involve is this. It would involve our maintaining that, where the total consequences of any actual voluntary action have more intrinsic value than those of the possible alternatives, it *absolutely always* happens to be true that they *also* contain more pleasure, although, in other cases, we know that degree of intrinsic value is by no means always in proportion to quantity of pleasure contained. And, of course, it is theoretically possible that this should be so: it is *possible* that the total consequences of actual voluntary actions should form a complete exception to the general rule: that, in their case, what has more intrinsic value should *absolutely always* also contain more pleasure, although, in other cases, this is by no means always true: but anybody can see, I think, that, in the absence of strict proof that it is so, the probabilities are all the other way. It is, indeed, so far as I can see, quite impossible absolutely to *prove* either that it is so or that it is not so; because *actual* actions in this world are liable to

have such an immense number of indirect and remote consequences, which we cannot trace, that it is impossible to be quite certain how the *total* consequences of any two actions will compare either in respect of intrinsic value, or in respect of the quantity of pleasure they contain. It *may*, therefore, *possibly* be the case that quantity of pleasure *is*, as a matter of fact, a correct *criterion* of right and wrong, even if intrinsic value is *not* always in proportion to quantity of pleasure contained. But it is impossible to *prove* that it is a correct criterion, except by assuming that intrinsic value always *is* in proportion to quantity of pleasure. And most of those who have held the former view have, I think, in fact made this assumption, even if they have not definitely realized that they were making it.

Is this assumption true, then? Is it true that one whole will be intrinsically better than another, whenever and only when it contains more pleasure, no matter what the two may be like in other respects? It seems to me almost impossible that any one, who fully realizes the consequences of such a view, can possibly hold that it *is* true. It involves our saying, for instance, that a world in which absolutely nothing except pleasure existed—no knowledge, no love, no enjoyment of beauty, no moral qualities—must yet be intrinsically better—better worth creating—provided only the total quantity of pleasure in it were the least bit greater, than one in which all these things existed *as well as* pleasure. It involves our saying that, even if the total quantity of pleasure in each was exactly equal, yet the fact that all the beings in the one possessed in addition knowledge of many different kinds and a full appreciation of all that was beautiful or worthy of love in their world, whereas *none* of the beings in the other possessed any of these things, would give us no

reason whatever for preferring the former to the latter. It involves our saying that, for instance, the state of mind of a drunkard, when he is intensely pleased with breaking crockery, is just as valuable, in itself—just as well worth having, as that of a man who is fully realizing all that is exquisite in the tragedy of King Lear, provided only the mere quantity of pleasure in both cases is the same. Such instances might be multiplied indefinitely, and it seems to me that they constitute a *reductio ad absurdum* of the view that intrinsic value is always in proportion to quantity of pleasure. Of course, here again, the question is quite incapable of proof either way. And if anybody, after clearly considering the issue, does come to the conclusion that no one kind of enjoyment is ever intrinsically better than another, provided only that the pleasure in both is equally intense, and that, if we *could* get as much pleasure in the world, without needing to have any knowledge, or any moral qualities, or any sense of beauty, as we can get *with* them, then all these things would be entirely superfluous, there is no way of proving that he is wrong. But it seems to me almost impossible that anybody, who does really get the question clear, should take such a view; and, if anybody were to, I think it is self-evident that he would be wrong.

It may, however, be asked: If the matter is as plain as this, how has it come about that anybody ever has adopted the view that intrinsic value *is* always in proportion to quantity of pleasure, or has ever argued, as if it were so? And I think one chief answer to this question is that those who have done so have *not* clearly realized all the consequences of their view, partly because they have been too exclusively occupied with the particular question as to whether, in the case

of *the total consequences* of *actual* voluntary actions, degree of intrinsic value is not always in proportion to quantity of pleasure—a question which, as has been admitted, is, in itself, much more obscure. But there is, I think, another reason, which is worth mentioning, because it introduces us to a principle of great importance. It may, in fact, be held, with great plausibility, that no whole can ever have any intrinsic value *unless* it contains some pleasure; and it might be thought, at first sight, that this reasonable, and perhaps true, view could not possibly lead to the wholly unreasonable one that intrinsic value is always *in proportion* to quantity of pleasure: it might seem obvious that to say that nothing can be valuable *without* pleasure is a very different thing from saying that intrinsic value is always *in proportion* to pleasure. And it is, I think, in fact true that the two views are really as different as they seem, and that the latter does not at all follow from the former. But, if we look a little closer, we may, I think, see a reason why the latter should very naturally have been *thought* to follow from the former.

The reason is as follows. If we say that no whole can ever be intrinsically good, *unless* it contains some pleasure, we are, of course, saying that if from any whole, which is intrinsically good, we were to subtract all the pleasure it contains, the remainder, whatever it might be, would have no intrinsic goodness at all, but must always be either intrinsically *bad*, or else intrinsically indifferent: and this (if we remember our definition of intrinsic value) is the same thing as to say that this remainder actually *has* no intrinsic goodness at all, but always *is* either positively bad or indifferent. Let us call the pleasure which such a whole contains, A, and the whole remainder, whatever it may be, B. We are then saying that the whole A + B is intrinsically good,

but that B is *not* intrinsically good at all. Surely it seems to follow that the intrinsic value of A + B cannot possibly be greater than that of A by itself? How, it may be asked, could it possibly be otherwise? How, by adding to A something, namely B, which has *no* intrinsic goodness at all, could we possibly get a whole which has *more* intrinsic value than A? It may naturally seem to be self-evident that we could not. But, if so, then it absolutely follows that we can never increase the value of any whole whatever except by adding *pleasure* to it: we may, of course, *lessen* its value, by adding other things, e.g. by adding pain; but we can never *increase* it except by adding pleasure.

Now from this it does not, of course, follow strictly that the intrinsic value of a whole is always *in proportion* to the quantity of pleasure it contains in the special sense in which we have throughout been using this expression—that is to say, as meaning that it is in proportion to the *excess* of pleasure over pain, in one of the five senses explained in Chapter I. But it is surely very natural to think that it does. And it *does* follow that we must be wrong in the reasons we gave for disputing this proposition. It does follow that we must be wrong in thinking that by adding such things as knowledge or a sense of beauty to a world which contained a certain amount of pleasure, without adding any more pleasure, we could increase the intrinsic value of that world. If, therefore, we are to dispute the proposition that intrinsic value *is* always in proportion to quantity of pleasure we must dispute this argument. But the argument may seem to be almost indisputable. It has, in fact, been used as an argument in favour of the proposition that intrinsic value *is* always in proportion to quantity of pleasure,

and I think it has probably had much influence in inducing people to adopt that view, even if they have not expressly put it in this form.

How, then, can we dispute this argument? We might, of course, do so, by rejecting the proposition that no whole can ever be intrinsically good, *unless* it contains some pleasure; but, for my part, though I don't feel certain that this proposition *is* true, I also don't feel at all certain that it is *not* true. The part of the argument which it seems to me certainly can and ought to be disputed is another part—namely, the assumption that, where a whole contains two factors, A and B, and one of these, B, has no intrinsic goodness at all, the intrinsic value of the whole cannot be *greater* than that of the other factor, A. This assumption, I think, obviously rests on a still more general assumption, of which it is only a special case. The general assumption is: That where a whole consists of two factors A and B, the amount by which its intrinsic value exceeds that of one of these two factors must always be equal to that of the other factor. Our special case will follow from this general assumption: because it will follow that if B be intrinsically *indifferent*, that is to say, if its intrinsic value = 0, then the amount by which the value of the whole A + B exceeds the value of A must also = 0, that is to say, the value of the whole must be precisely *equal* to that of A; while if B be intrinsically *bad*, that is to say, if its intrinsic value is less than 0, then the amount by which the value of A + B will exceed that of A will also be less than 0, that is to say, the value of the whole will be *less* than that of A. Our special case does then follow from the general assumption; and nobody, I think, would maintain that the special case was true without maintaining that the general assumption was also true. The

general assumption may, indeed, very naturally seem to be self-evident: it has, I think, been generally assumed that it is so: and it may seem to be a mere deduction from the laws of arithmetic. But, so far as I can see, it is *not* a mere deduction from the laws of arithmetic, and, so far from being self-evident, is certainly untrue.

Let us see exactly what we are saying, if we deny it. We are saying that the fact that A and B *both* exist together, together with the fact that they have to one another any relation which they do happen to have (when they exist together, they always must have *some* relation to one another; and the precise nature of the relation certainly may in some cases make a great difference to the value of the whole state of things, though, perhaps, it need not in all cases)—that these two facts *together* must have a certain amount of intrinsic value, that is to say must be either intrinsically good, or intrinsically bad, or intrinsically indifferent, and that the amount by which this value exceeds the value which the existence of A would have, if A existed quite alone, *need* not be equal to the value which the existence of B would have, if B existed quite alone. This is all that we are saying. And can any one pretend that such a view necessarily contradicts the laws of arithmetic? or that it is self-evident that it cannot be true? I cannot see any ground for saying so; and if there is no ground, then the argument which sought to show that we can never add to the value of any whole *except* by adding pleasure to it, is entirely baseless.

If, therefore, we reject the theory that intrinsic value is always in proportion to quantity of pleasure, it does seem as if we may be compelled to accept the principle that *the amount by which the value of a whole exceeds that of one of its factors is not necessarily equal*

to that of the remaining factor—a principle which, if true, is very important in many other cases. But, though at first sight this principle may seem paradoxical, there seems to be no reason why we should not accept it; while there are other independent reasons why we should accept it. And, in any case, it seems quite clear that the degree of intrinsic value of a whole is *not* always in proportion to the quantity of pleasure it contains.

But, if we do reject this theory, what, it may be asked, can we substitute for it? How can we answer the question, what kinds of consequences are intrinsically better or worse than others?

We may, I think, say, first of all, that for the same reason for which we have rejected the view that intrinsic value is always in proportion to quantity of pleasure, we must also reject the view that it is always in proportion to the quantity of any other *single* factor whatever. Whatever single kind of thing may be proposed as a measure of intrinsic value, instead of pleasure—whether knowledge, or virtue, or wisdom, or love—it is, I think, quite plain that it is not such a measure; because it is quite plain that, however valuable any one of these things may be, we may always add to the value of a whole which contains any one of them, not only by adding more of that one, but also *by adding something else instead.* Indeed, so far as I can see, there is no characteristic whatever which always distinguishes every whole which has greater intrinsic value from every whole which has less, *except* the fundamental one that it would always be the duty of every agent to prefer the better to the worse, if he had to choose between a pair of actions, of which they would be the *sole* effects. And similarly, so far as I can see, there is no characteristic whatever which

belongs to all things that are intrinsically *good* and only to them—except simply the one that they all *are* intrinsically good and *ought* always to be preferred to *nothing at all*, it we had to choose between an action whose sole effect would be one of them and one which would have no effects whatever. The fact is that the view which seems to me to be true is the one which, apart from theories, I think every one would naturally take, namely, that there are an *immense variety* of different things, *all* of which are intrinsically good; and that though all these things may perhaps have some characteristic *in common*, their variety is so great that they have none, which, *besides* being common to them all, is also *peculiar* to them—that is to say, which never belongs to anything which is intrinsically bad or indifferent. All that can, I think, be done by way of making plain what kinds of things are intrinsically good or bad, and what are better or worse than others, is to classify some of the chief kinds of each, pointing out what the factors are upon which their goodness or badness depends. And I think this is one of the most profitable things which can be done in Ethics, and one which has been too much neglected hitherto. But I have not space to attempt it here.

I have only space for two final remarks. The first is that there do seem to be two important characteristics, which are *common* to absolutely all intrinsic goods, though not peculiar to them. Namely (1) it does seem as if nothing can be an intrinsic good unless it contains *both* some feeling and *also* some other form of consciousness; and, as we have said before, it seems possible that amongst the feelings contained must always be some amount of pleasure. And (2) it does also seem as if every intrinsic good must be a complex whole containing a considerable variety of different

factors — as if, for instance, nothing so simple as pleasure by itself, however intense, could ever be any good. But it is important to insist (though it is obvious) that neither of these characteristics is *peculiar* to intrinsic goods: they may obviously *also* belong to things bad and indifferent. Indeed, as regards the first, it is not only true that many wholes which contain both feeling and some other form of consciousness are intrinsically bad; but it seems also to be true that nothing can be intrinsically bad, *unless* it contains some feeling.

The other final remark is that we must be very careful to distinguish the two questions (1) whether, and in what degree, a thing is *intrinsically* good and bad, and (2) whether, and in what degree, it is capable of adding to or subtracting from the intrinsic value of a whole of which it forms a part, from a third, entirely different question, namely (3) whether, and in what degree, a thing is *useful* and has good *effects*, or *harmful* and has *bad* effects. All three questions are very liable to be confused, because, in common life, we apply the names 'good' and 'bad' to things of all three kinds indifferently: when we say that a thing is 'good' we may mean either (1) that it is intrinsically good or (2) that it adds to the value of many intrinsically good wholes or (3) that it is useful or has good effects; and similarly when we say that a thing is bad we may mean any one of the three corresponding things. And such confusion is very liable to lead to mistakes, of which the following are, I think, the commonest. In the first place, people are apt to assume with regard to things, which really are very good indeed in senses (1) or (2), that they are scarcely any good at all, simply because they do not seem to be of much *use*—that is to say, to lead to *further* good effects; and similarly, with regard

to things which really are very bad in senses (1) or (2), it is very commonly assumed that there cannot be much, if any, harm in them, simply because they do not seem to lead to *further* bad results. Nothing is commoner than to find people asking of a good thing: What *use* is it? and concluding that, if it is no use, it cannot be any good; or asking of a bad thing: What harm does it do? and concluding that if it *does* no harm, there cannot be any harm *in* it. Or, again, by a converse mistake, of things which really are very useful, but are not good at all in senses (1) and (2), it is very commonly assumed that they *must* be good in one or both of these two senses. Or again, of things, which really are very good in senses (1) and (2), it is assumed that, because they are good, they cannot possibly do harm. Or finally, of things, which are neither intrinsically good nor useful, it is assumed that they cannot be any good at all, although in fact they are very good in sense (2). All these mistakes are liable to occur, because, in fact, the degree of goodness or badness of a thing in any one of these three senses is by no means always in proportion to the degree of its goodness or badness in either of the other two; but if we are careful to distinguish the three different questions, they can, I think, all be avoided.

NOTE ON BOOKS

If the reader wishes to form an impartial judgement as to what the fundamental problems of Ethics really are, and what is the true answer to them, it is of the first importance that he should not confine himself to reading works of any one single type, but should realize what extremely different sorts of things have seemed to different writers, of acknowledged reputation, to be the most important things to be said about the subject. For this purpose he should, I think, read, if possible, and compare with one another, *all* of the following works:

1. Some of the dialogues of Plato (translated by Jowett). Among the shorter dialogues the *Protagoras*, the *Gorgias*, and the *Philebus* deal almost exclusively with fundamental ethical questions, and may be taken as typical examples of Plato's method of dealing with Ethics; but the reader should, if possible, read also the whole of the *Republic*, because, though, in the main, it is concerned with points of comparative detail, it contains, in various places, discussions which are of great importance for understanding Plato's general view.
2. Aristotle's *Nicomachean Ethics*. (There are several English translations.)
3. Hume's *Enquiry concerning the Principles of Morals*.
4. Kant's *Fundamental Principles of the Metaphysic of Morals*. (Translated, along with other works, under the title *Kant's Theory of Ethics*, by T. K. Abbott: Longmans, Green & Co.)
5. John Stuart Mill's *Utilitarianism*.
6. Henry Sidgwick's *Methods of Ethics* (Macmillan & Co.).
7. Herbert Spencer's *Data of Ethics* (forming the first part of his two volumes on *The Principles of Ethics*, but also published separately).
8. T. H. Green's *Prolegomena to Ethics* (Clarendon Press).

I have selected these works as being enough, but not more than enough, to give a sufficient idea of the extremely different way in which writers, who are still considered by many people to be among the best worth reading on the subject, have dealt with it. No doubt, in some cases, other works, equally well worth reading, and equally typical of the sort of differences I want to emphasize, might be substituted for some of those I have mentioned; but these are, I think, *as* good as any for the purposes of illustration, and hardly one of them could be omitted without serious loss, unless some

other work, typical of the same method of treatment, were substituted for it.

For guidance in his further reading, so far as writers no longer living are concerned, the reader may be referred to Sidgwick's *Outlines of the History of Ethics* (Macmillan & Co.), from which he will be able to judge what other writers it is likely to be most profitable for him to study, and which is also well worth reading on its own account. And, if he wishes to become acquainted with the principal works on Ethics which have been written by writers still living, I think I can hardly do better than recommend him to read, first of all, Dr. Hastings Rashdall's *Theory of Good and Evil* (Clarendon Press, 1907). This book will, I think, give a fair idea of the sort of questions which are still being discussed at the present day, and it also contains references to the most important works of other living writers, sufficient to enable the reader to make his own choice of further reading.

For further explanation of the views advocated in the present work the reader may be referred to the author's *Principia Ethica* (Cambridge University Press, 1903), which presents the same general view in a rather different form, and which also contains discussions on various points entirely omitted here from lack of space.