

CHAPTER VI FREE WILL

THROUGHOUT the last three chapters we have been considering various objections which might be urged against the theory stated in Chapters I and II. And the very last objection which we considered was one which consisted in asserting that the question whether an action is right or wrong does *not* depend upon its *actual* consequences, because whenever the consequences, *so far as the agent can foresee*, are *likely* to be the best possible, the action is always right, even if they are not *actually* the best possible. In other words, this objection rested on the view that right and wrong depend, in a sense, upon what the agent *can know*. And in the present chapter I propose to consider objections, which rest, instead of this, upon the view that right and wrong depend upon what the agent *can do*.

Now it must be remembered that, *in a sense*, our original theory does hold and even insists that this is the case. We have, for instance, frequently referred to it in the last chapter as holding that an action is only right, if it produces the best *possible* consequences; and by 'the best *possible* consequences' was meant 'consequences at least as good as would have followed from any action which the agent *could* have done instead'. It does, therefore, hold that the question whether an action is right or wrong does always depend upon a comparison of its consequences with those of all the other actions which the agent *could* have done instead. It assumes, therefore, that wherever a voluntary action is right or wrong (and we have throughout only been

talking of *voluntary* actions), it is true that the agent *could*, in a sense, have done something else instead. This is an absolutely essential part of the theory.

But the reader must now be reminded that all along we have been using the words 'can', 'could', and 'possible' *in a special sense*. It was explained in Chapter I (pp. 20-21), that we proposed, purely for the sake of brevity, to say that an agent *could* have done a given action, which he didn't do, wherever it is true that he could have done it, *if* he had chosen; and similarly by what he *can* do, or what is *possible*, we have always meant merely what is possible, *if* he chooses. Our theory, therefore, has not been maintaining, after all, that right and wrong depend upon what the agent absolutely *can* do, but only on what he can do, *if* he chooses. And this makes an immense difference. For, by confining itself in this way, our theory avoids a controversy, which cannot be avoided by those who assert that right and wrong depend upon what the agent absolutely *can* do. There are few, if any, people who will *expressly* deny that we very often really could, *if* we had chosen, have done something different from what we actually did do. But the moment it is asserted that any man ever absolutely *could* have done anything other than what he did do, there are many people who *would* deny this. The view, therefore, which we are to consider in this chapter—the view that right and wrong depend upon what the agent absolutely *can* do—at once involves us in an extremely difficult controversy—the controversy concerning Free Will. There are many people who strenuously deny that any man ever *could* have done anything other than what he actually did do, or ever *can* do anything other than what he *will* do; and there are others who assert the opposite equally strenuously. And whichever view be

held is, if *combined* with the view that right and wrong depend upon what the agent absolutely *can* do, liable to contradict our theory very seriously. Those who hold that no man ever *could* have done anything other than what he did do, are, if they *also* hold that right and wrong depend upon what we *can* do, logically bound to hold that no action of ours is ever right and none is ever wrong; and this is a view which is, I think, often actually held, and which, of course, constitutes an extremely serious and fundamental objection to our theory: since our theory implies, on the contrary, that we very often do act *wrongly*, if never quite rightly. Those, on the other hand, who hold that we absolutely *can* do things, which we don't do, and that right and wrong depend upon what we thus *can* do, are also liable to be led to contradict our theory, though for a different reason. Our theory holds that, provided a man could have done something else, *if* he had chosen, that is sufficient to entitle us to say that his action really is either right or wrong. But those who hold the view we are considering will be liable to reply that this is by no means sufficient: that to say that it *is* sufficient, is entirely to misconceive the nature of right and wrong. They will say that, in order that an action may be *really* either right or wrong, it is absolutely essential that the agent should have been *really able* to act differently, able in some sense quite other than that of merely being able, *if* he had chosen. *If* all that were really ever true of us were merely that we could have acted differently, *if* we had chosen, then, these people would say, it really would be true that none of our actions are ever right and that none are ever wrong. They will say, therefore, that our theory entirely misses out one absolutely essential condition of right and wrong—the condition that, for

an action to be right or wrong, it must be *freely* done. And moreover, many of them will hold also that the class of actions which we absolutely *can* do is often not identical with those which we can do, *if* we choose. They may say, for instance, that very often an action, which we *could* have done, *if* we had chosen, is nevertheless an action which we *could not* have done; and that an action is always right, if it produces as good consequences as any other action which we really *could* have done instead. From which it will follow that many actions which our theory declares to be *wrong*, will, according to them, be right, because these actions really are the best of all that we *could* have done, though *not* the best of all that we could have done, *if* we had chosen.

Now these objections seem to me to be the most serious which we have yet had to consider. They seem to me to be serious because (1) it is very difficult to be sure that right and wrong do not really depend, as they assert, upon what we *can* do and not merely on what we can do, *if* we choose; and because (2) it is very difficult to be sure in what sense it is true that we ever *could* have done anything different from what we actually did do. I do not profess to be sure about either of these points. And all that I can hope to do is to point out certain facts which do seem to me to be clear, though they are often overlooked; and thus to isolate clearly for the reader's decision, those questions which seem to me to be really doubtful and difficult.

Let us begin with the question: Is it ever true that a man *could* have done anything else, except what he actually did do? And, first of all, I think I had better explain exactly how this question seems to me to be related to the question of Free Will. For it is a fact that, in many discussions about Free Will, this precise

question is never mentioned at all; so that it might be thought that the two have really nothing whatever to do with one another. And indeed some philosophers do, I think, definitely imply that they *have* nothing to do with one another: they seem to hold that our wills can properly be said to be free even if we *never* can, in any sense at all, do anything else except what, in the end, we actually do do. But this view, if it is held, seems to me to be plainly a mere abuse of language. The statement that we have Free Will is certainly ordinarily understood to imply that we really sometimes have the power of acting differently from the way in which we actually do act; and hence, if anybody tells us that we have Free Will, while at the same time he means to deny that we ever have such a power, he is simply misleading us. We certainly have *not* got Free Will, in the ordinary sense of the word, if we never really *could*, in any sense at all, have done anything else than what we did do; so that, in this respect, the two questions certainly are connected. But, on the other hand, the mere fact (if it is a fact) that we sometimes *can*, in *some* sense, do what we don't do, does not necessarily entitle us to say that we *have* Free Will. We certainly *haven't* got it, *unless* we can; but it doesn't follow that we *have* got it, even if we *can*. Whether we have or not will depend upon the precise sense in which it is true that we can. So that even if we do decide that we really *can* often, in *some* sense, do what we don't do, this decision by itself does not entitle us to say that we have Free Will.

And the first point about which we can and should be quite clear is, I think, this: namely, that we certainly often *can*, in *some* sense, do what we don't do. It is, I think, quite clear that this is so; and also very important that we should realize that it is so. For many

people are inclined to assert, quite without qualification: No man ever *could*, on any occasion, have done anything else than what he actually did do on that occasion. By asserting this quite simply, without qualification, they imply, of course (even if they do not mean to imply), that there is *no* proper sense of the word 'could', in which it is true that a man *could* have acted differently. And it is this implication which is, I think, quite certainly absolutely false. For this reason, anybody who asserts, without qualification, 'Nothing ever *could* have happened, except what actually did happen', is making an assertion which is quite unjustifiable, and which he himself cannot help constantly contradicting. And it is important to insist on this, because many people do make this unqualified assertion, without seeing how violently it contradicts what they themselves, and all of us, believe, and rightly believe, at other times. If, indeed, they insert a qualification—if they merely say. 'In *one* sense of the word "*could*" nothing ever *could* have happened, except what did happen', then, they may perhaps be perfectly right: we are not disputing that they may. All that we are maintaining is that, in *one* perfectly proper and legitimate sense of the word 'could', and that one of the very commonest senses in which it is used, it is quite certain that some things which didn't happen *could* have happened. And the proof that this is so, is simply as follows.

It is impossible to exaggerate the frequency of the occasions on which we *all* of us make a distinction between two things, neither of which *did* happen—a distinction which we express by saying, that whereas the one *could* have happened, and other could *not*. No distinction is commoner than this. And no one, I think, who fairly examines the instances in which we

make it, can doubt about three things: namely (1) that very often there really is *some* distinction between the two things, corresponding to the language which we use; (2) that this distinction, which really *does* subsist between the things, is *the* one which we mean to express by saying that the one was possible and the other impossible; and (3) that this way of expressing it is a perfectly proper and legitimate way. But if so, it absolutely follows that one of the commonest and most legitimate usages of the phrases 'could' and 'could not' is to express a difference, which often really does hold between two things *neither* of which did actually happen. Only a few instances need be given. I *could* have walked a mile in twenty minutes this morning, but I certainly *could not* have run two miles in five minutes. I did not, *in fact*, do either of these two things; but it is pure nonsense to say that the mere fact that I *did* not, does away with the distinction between them, which I express by saying that the one *was* within my powers, whereas the other was *not*. *Although* I did neither, yet the one was certainly *possible* to me in a sense in which the other was totally *impossible*. Or, to take another instance: It is true, as a rule, that cats *can* climb trees, whereas dogs *can't*. Suppose that on a particular afternoon neither A's cat nor B's dog *do* climb a tree. It is quite absurd to say that this mere fact proves that we must be wrong if we say (as we certainly often should say) that the cat *could* have climbed a tree, though she didn't, whereas the dog *couldn't*. Or, to take an instance which concerns an inanimate object. Some ships *can* steam 20 knots, whereas others *can't* steam more than 15. And the mere fact that, on a particular occasion, a 20-knot steamer *did* not *actually* run at this speed certainly does not entitle us to say that she *could* not have done

'so, in the sense in which a 15-knot one *could* not. On the contrary, we all can and should distinguish between cases in which (as, for instance, owing to an accident to her propeller) she did not, *because* she could not, and cases in which she did not, *although* she *could*. Instances of this sort might be multiplied quite indefinitely; and it is surely quite plain that we all of us do *continually* use such language: we continually, when considering two events, neither of which *did* happen, distinguish between them by saying that whereas the one *was* possible, though it didn't happen, the other was *impossible*. And it is surely quite plain that what we mean by this (whatever it may be) is something which is often perfectly true. But, if so, then anybody who asserts, without qualification, 'Nothing ever *could* have happened, except what did happen', is simply asserting what is false.

It is, therefore, quite certain that we often *could* (in *some* sense) have done what we did not do. And now let us see how this fact is related to the argument by which people try to persuade us that it is *not* a fact.

The argument is well known: it is simply this. It is assumed (for reasons which I need not discuss) that absolutely everything that happens has a *cause* in what precedes it. But to say this is to say that it follows *necessarily* from something that preceded it; or, in other words, that, once the preceding events which are its cause had happened, it was absolutely *bound* to happen. But to say that it was *bound* to happen, is to say that nothing else *could* have happened instead; so that, if *everything* has a cause, *nothing* ever could have happened except what did happen.

And now let us assume that the premise of this argument is correct: that everything really *has* a cause. What really follows from it? Obviously all that

follows is that, in *one* sense of the word 'could', nothing ever *could* have happened, except what did happen. This really *does* follow. But, *if* the word 'could' is ambiguous—if, that is to say, it is used in different senses on different occasions—it is obviously quite possible that though, in *one* sense, nothing ever could have happened except what did happen, yet in *another* sense, it may at the same time be perfectly true that some things which did not happen *could* have happened. And can anybody undertake to assert with certainty that the word 'could' is *not* ambiguous? that it may not have more than one legitimate sense? *Possibly* it is not ambiguous; and, *if* it is not, then the fact that some things, which did not happen, *could* have happened, really would contradict the principle that everything has a cause; and, in that case, we should, I think, have to give up this principle, because the fact that we often *could* have done what we did not do, is so certain. But the assumption that the word 'could' is *not* ambiguous is an assumption which certainly should not be made without the clearest proof. And yet I think it often is made, without any proof at all; simply because it does not occur to people that words often are ambiguous. It is, for instance, often assumed, in the Free Will controversy, that the question at issue is solely as to whether everything is caused, or whether acts of will are sometimes uncaused. Those who hold that we *have* Free Will, think themselves bound to maintain that acts of will sometimes have *no* cause; and those who hold that everything is caused think that this proves completely that we have not Free Will. But, in fact, it is extremely doubtful whether Free Will is at all inconsistent with the principle that everything is caused. Whether it is or not, all depends on a very difficult question as to the meaning of the

word 'could'. All that is certain about the matter is (1) that, if we have Free Will, it must be true, in *some* sense, that we sometimes *could* have done, what we did not do; and (2) that, if everything is caused, it must be true, in *some* sense, that we *never could* have done, what we did not do. What is very *uncertain*, and what certainly needs to be investigated, is whether these two meanings of the word 'could' are the same.

Let us begin by asking: What is the sense of the word 'could', in which it is so certain that we often *could* have done, what we did not do? What, for instance, is the sense in which I *could* have walked a mile in twenty minutes this morning, though I did not? There is one suggestion, which is very obvious: namely, that what I mean is simply after all that I could, *if* I had chosen; or (to avoid a possible complication) perhaps we had better say 'that I *should*, *if* I had chosen'. In other words, the suggestion is that we often use the phrase '*I could*' simply and solely as a short way of saying '*I should*, if I had chosen'. And in all cases, where it is certainly true that we *could* have done, what we did not do, it is, I think, very difficult to be quite sure that this (or something similar) is *not* what we mean by the word 'could'. The case of the ship may seem to be an exception, because it is certainly not true that she would have steamed twenty knots if *she* had chosen; but even here it seems possible that what we mean is simply that she *would*, *if the men on board of her* had chosen. There are certainly good reasons for thinking that we *very often* mean by 'could' merely 'would, *if* so and so had chosen'. And if so, when we have a sense of the word 'could' in which the fact that we often *could* have done what we did not do, is perfectly compatible with the principle that everything has a cause: for to say

that, *if* I had performed a certain act of will, I should have done something which I did not do, in no way contradicts this principle.

And an additional reason for supposing that this *is* what we often mean by 'could', and one which is also a reason why it is important to insist on the obvious fact that we very often really *should* have acted differently, *if* we had willed differently, is that those who deny that we ever *could* have done anything, which we did not do, often speak and think as if this really did involve the conclusion that we never should have acted differently, even *if* we had willed differently. This occurs, I think, in two chief instances—one in reference to the future, the other in reference to the past. The first occurs when, because they hold that nothing *can* happen, except what *will* happen, people are led to adopt the view called Fatalism—the view that *whatever we will*, the result will always be the same; that it is, therefore, *never* any use to make one choice rather than another. And this conclusion will really follow if by 'can' we mean '*would* happen, even *if* we were to will it'. But it is certainly untrue, and it certainly does not follow from the principle of causality. On the contrary, reasons of exactly the same sort and exactly as strong as those which lead us to suppose that everything has a cause, lead to the conclusion that if we choose one course, the result will *always* be different in *some* respect from what it would have been, if we had chosen another; and we know also that the difference would *sometimes* consist in the fact that *what* we chose would come to pass. It is certainly often true of the future, therefore, that whichever of two actions we *were* to choose, *would* actually be done, although it is quite certain that only one of the two *will* be done.

And the second instance, in which people are apt to speak and think, as if, *because* no man ever *could* have done anything but what he did do, it follows that he would not, even *if* he had chosen, is as follows. Many people seem, in fact, to conclude directly from the first of these two propositions, that we can never be justified in praising or blaming a man for anything that he does, or indeed for making any distinction between what is right or wrong, on the one hand, and what is lucky or unfortunate on the other. They conclude, for instance, that there is never any reason to treat or to regard the voluntary commission of a crime in any different way from that in which we treat or regard the involuntary catching of a disease. The man who committed the crime *could* not, they say, have helped committing it any more than the other man could have helped catching the disease; both events were equally inevitable; and though both may have very bad consequences and equally bad ones—there is no justification whatever, they say, for the distinction we make between them when we say that the commission of the crime was *wrong*, or that the man was morally to blame for it, whereas the catching of the disease was *not* wrong and the man was not to blame for it. And this conclusion, again, will really follow if by '*could* not' we mean '*would* not, even if he had willed to avoid it'. But the point I want to make is, that it follows *only* if we make this assumption. That *is* to say, the mere fact that the man *would* have succeeded in avoiding the crime, *if* he had chosen (which *is* certainly often true), whereas the other man would *not* have succeeded in avoiding the disease, *even* if he had chosen (which is certainly also often true) gives an ample justification for regarding and treating the two

cases differently. It gives such a justification, because, where the occurrence of an event *did* depend upon the will, there, by acting on the will (as we may do by blame or punishment) we have often a reasonable chance of preventing similar events from recurring in the future; whereas, where it did *not* depend upon the will, we have no such chance. We may, therefore, fairly say that those who speak and think, as if a man who brings about a misfortune *voluntarily* ought to be treated and regarded in exactly the same way as one who brings about an equally great misfortune *involuntarily*, are speaking and thinking *as if* it were not true that we ever should have acted differently, even *if* we had willed to do so. And that is why it is extremely important to insist on the absolute certainty of the fact that we often really *should* have acted differently, *if* we had willed differently.

There is, therefore, much reason to think that when we say that we *could* have done a thing which we did not do, we *often* mean merely that we *should* have done it, *if* we had chosen. And if so, then it is quite certain that, in *this* sense, we often really *could* have done what we did not do, and that this fact is in no way inconsistent with the principle that everything has a cause. And for my part I must confess that I cannot feel certain that this may not be *all* that we usually mean and understand by the assertion that we have Free Will; so that those who deny that we have it are really denying (though, no doubt, often unconsciously) that we ever *should* have acted differently, even if we had willed differently. It has been sometimes held that this *is* what we mean; and I cannot find any conclusive argument to the contrary. And if it is *have* we mean, then it absolutely follows that we really have Free Will, and also that this fact is quite consistent

with the principle that everything has a cause; and it follows also that our theory will be perfectly right, when it makes right and wrong depend on what we *could* have done, *if* we had chosen.

But, no doubt, there are many people who will say that this is *not* sufficient to entitle us to say that we have Free Will; and they will say this for a reason, which certainly has some plausibility, though I cannot satisfy myself that it is conclusive. They will say, namely: Granted that we often *should* have acted differently, *if* we had chosen differently, yet it is not true that we have Free Will, unless it is *also* often true in such cases that we *could* have *chosen* differently. The question of Free Will has been thus represented as being merely the question whether we ever *could* have chosen, what we did not choose, or ever *can* choose, what, in fact, we shall not choose. And since there is some plausibility in this contention, it is, I think, worth while to point out that here again it is absolutely certain that, in two different senses, at least, we often *could* have chosen, what, in fact, we did not choose; and that in neither sense does this fact contradict the principle of causality.

The first is simply the old sense over again. If by saying that we *could* have done, what we did not do, we often mean merely that we *should* have done it, *if* we had chosen to do it, then obviously, by saying that we *could* have *chosen* to do it, we may mean merely that we *should* have so chosen, *if* we had chosen *to make the choice*. And I think there is no doubt it is often true that we should have chosen to do a particular thing *if* we had chosen to make the choice; and that this is a very important sense in which it is often in our power to make a choice. There certainly is such a thing as making an effort to induce ourselves to *choose*

a particular course; and I think there is no doubt that often if we *had* made such an effort, we *should* have made a choice, which we did not in fact make.

And besides this, there is another sense in which, whenever we have several different courses of action in view, it is *possible* for us to choose any one of them; and a sense which is certainly of some practical importance, even if it goes no way to justify us in saying that we have Free Will. This sense arises from the fact that in such cases we can hardly ever *know for certain* beforehand, *which* choice we actually *shall* make; and one of the commonest senses of the word 'possible' is that in which we call an event 'possible' when no man can *know for certain* that it will *not* happen. It follows that almost, if not quite always, when we make a choice, after considering alternatives, it *was* possible that we should have chosen one of these alternatives, which we did not actually choose; and often, of course, it was not only possible, but highly probable, that we should have done so. And this fact is certainly of practical importance, because many people are apt much too easily to assume that it is quite certain that they *will not* make a given choice, which they know they ought to make, if it were possible; and their belief that they *will not* make it tends, of course, to prevent them from making it. For this reason it is important to insist that they can hardly ever know for certain with regard to any given choice that they will *not* make it.

It is, therefore, quite certain (1) that we often *should* have *acted* differently, if we had chosen to; (2) that similarly we often should have *chosen* differently, if we had chosen so to choose; and (3) that it was almost always *possible* that we should have chosen differently, in the sense that no man could know for certain that

we should *not* so choose. All these three things are facts, and all of them are quite consistent with the principle of causality. Can anybody undertake to say for certain that none of these three facts and *no* combination of them will justify us in saying that we have Free Will? Or, suppose it granted that we have not Free Will, unless it is often true that we *could* have chosen, what we did not choose:—Can any defender of Free Will, or any opponent of it, show conclusively that what he means by '*could* have chosen' in this proposition, is anything different from the two certain facts, which I have numbered (2) and (3), or some combination of the two? Many people, no doubt, will still insist that these two facts alone are by no means sufficient to entitle us to say that we have Free Will: that it must be true that we were *able* to choose, in some quite other sense. But nobody, so far as I know, has ever been able to tell us exactly what that sense is. For my part, I can find no conclusive argument to show either that some such other sense of 'can' is necessary, or that it is not. And, therefore, this chapter must conclude with a doubt. It is, I think, possible that, instead of saying, as our theory said, that an action is only right, when it produces consequences as good as any which would have followed from any other action which the agent *would* have done, *if* he had chosen, we should say instead that it is right whenever and only when the agent *could not have done* anything which would have produced better consequences: and that this '*could not have done*' is *not* equivalent to '*would not have done, if he had chosen*', but is to be understood in the sense, whatever it may be, which is sufficient to entitle us to say that we have Free Will. If so, then our theory would be wrong, just to this extent.