

Notice the difference here from the gravitational case. Our inner experiences are not, really, up and down from one another. No harm then in dissociating our inner experience of down from the real down-ness relation in the world and then identifying the latter with a relation characterized in terms of the gradient of the gravitational potential. But inner events are *really* later and earlier than one another and our 'pretty simple, immediate experience' of this relation cannot with impunity be detached as merely a causally induced secondary quality not properly thought of as a direct experience of the real afterness relation which exists in the world only as an entropically characterizable relation.

The following quote from Eddington suggests that it is something like this argument which is at the root of many of the strongly felt but not very well expressed objections to the plausibility of the entropic theory. It is important to note that this quote is from one of the earliest expositors of the entropic theory of time direction as I have described it.

In any attempt to bridge the domains of experience belonging to the spiritual and physical sides of our nature, Time occupies the key position. I have already referred to its dual entry into our consciousness—through the sense organs which relate it to the other entities of the physical world, and directly through a kind of private door into the mind. . . . Whilst the physicist would generally say that the matter of this familiar table is *really* a curvature of space, and its colour is really electromagnetic wavelength, I do not think he would say that the familiar moving on of time is *really* an entropy-gradient. . . . Our trouble is that we have to associate two things, both of which we more or less understand, and, so as we understand them, they are utterly different. It is absurd to pretend that we are in ignorance of the nature of organization in the external world in the same way that we are ignorant of the intrinsic nature of potential. It is absurd to pretend that we have no justifiable conception of 'becoming' in the external world. That dynamic quality—that significance which makes a development from future to past farcical—has to do much more than pull a trigger of a nerve. It is so welded into our consciousness that a moving on of time is a condition of consciousness. We have direct insight into 'becoming' which sweeps aside all symbolic knowledge as on an inferior plane. If I grasp the notion of existence because I myself exist, I grasp the notion of becoming because I myself become. It is the innermost Ego of all which *is* and *becomes*.¹¹

I don't pretend to understand all that Eddington is saying here, nor to be able to give a really coherent version of my own arguments above. I do think, however, that it is very clear that our ultimate view of the world will require a subtle and careful weaving together of the naturalistic reduction of science which proceeds by theoretical identification with the conceptual reduction of philosophy which proceeds by epistemic analysis. Until we have such a systematic overall account I think that the ultimate status of an entropic theory of time order will be in doubt.

¹¹ Eddington, *The Nature of the Physical World*, 91–7.

VII

BRINGING ABOUT THE PAST

MICHAEL DUMMETT

I observe first that there is a genuine sense in which the causal relation has a temporal direction: it is associated with the direction earlier-to-later rather than with the reverse. I shall not pause here to achieve a precise formulation of the sense in which this association holds; I think such a formulation can be given without too much difficulty, but it is not to my present purpose to do this. What I do want to assert is the following: so far as I can see, this association of causality with a particular temporal direction is not merely a matter of the way we speak of causes, but has a genuine basis in the way things happen. There is indeed an asymmetry in respect of past and future in the way in which we describe events when we are considering them as standing in causal relations to one another; but I am maintaining that this reflects an objective asymmetry in nature. I think that this asymmetry would reveal itself to us even if we were not *agents* but mere *observers*. It is indeed true, I believe, that our concept of cause is bound up with our concept of intentional action: if an event is properly said to cause the occurrence of a subsequent or simultaneous event, I think it necessarily follows that, if we can find any way of bringing about the former event (in particular, if it is itself a voluntary human action), then it must make sense to speak of bringing it about *in order that* the subsequent event should occur. Moreover, I believe that this connection between something's being a cause and the possibility of using it in order to bring about its effect plays an essential rôle in the fundamental account of how we ever come to accept causal laws: that is, that we could arrive at any causal beliefs only by beginning with those in which the cause is a voluntary action of ours. Nevertheless, I am inclined to think that we could have some kind of concept of cause, although one differing from that we now have, even if we were mere observers and not agents at all—a kind of intelligent tree. And I also think that even in this case the asymmetry of cause with respect to temporal direction would reveal itself to us.

Michael Dummett, 'Bringing About the Past'. First published in the *Philosophical Review*, 73 (1964): 338–59. The text as reprinted here includes some small changes made by the author when this essay was reprinted in his *Truth and Other Enigmas* (London: Duckworth, 1978).

To see this, imagine ourselves observing events in a world just like the actual one, except that the order of events is reversed. There are indeed enormous difficulties in describing such a world if we attempt to include human beings in it, or any other kind of creature to whom can be ascribed intention and purpose (there would also be a problem about memory). But, so far as I can see, there is no difficulty whatever if we include in this world only plants and inanimate objects. If we imagine ourselves as intelligent trees observing such a world and communicating with one another, but unable to intervene in the course of events, it is clear that we should have great difficulty in arriving at causal explanations that accounted for events in terms of the processes which had *led up to* them. The sapling grows gradually smaller, finally reducing itself to an apple pip; then an apple is gradually constituted around the pip from ingredients found in the soil; at a certain moment the apple rolls along the ground, gradually gaining momentum, bounces a few times, and then suddenly takes off vertically and attaches itself with a snap to the bough of an apple tree. Viewed from the standpoint of gross observation, this process contains many totally unpredictable elements: we cannot, for example, explain, by reference to the conditions obtaining at the moment when the apple started rolling, why it started rolling at that moment or in that direction. Rather, we should have to substitute a system of explanations of events in terms of the processes that led back to them from some subsequent moment. If through some extraordinary chance we, in this world, could consider events from the standpoint of the microscopic, the unpredictability would disappear theoretically ('in principle') although not in practice; but we should be left—so long as we continued to try to give causal explanations on the basis of what leads up to an event—with inexplicable coincidences. 'In principle' we could, by observing the movements of the molecules of the soil, predict that at a certain moment they were going to move in such a way as to combine to give a slight impetus to the apple, and that this impetus would be progressively reinforced by other molecules along a certain path, so as to cause the apple to accelerate in such a way that it would end up attached to the apple tree. But not only could we not make such predictions in practice: the fact that the 'random' movements of the molecules should happen to work out in such a way that all along the path the molecules always happened to be moving in the same direction at just the moment that the apple reached that point, and, above all, that these movements always worked in such a way as to leave the apple attached to an *apple* tree and not to any other tree or any other object—these facts would cry out for explanation, and we should be unable to provide it.

I should say, then, that, so far as the concept of cause possessed by mere observers rather than agents is concerned, the following two theses hold: (i)

the world is such as to make appropriate a notion of causality associated with the earlier-to-later temporal direction rather than its reverse; (ii) we can conceive of a world in which a notion of causality associated with the opposite direction would have been more appropriate and, so long as we consider ourselves as mere observers of such a world, there is no particular conceptual difficulty about the conception of such a backwards causation. There are, of course, regions of which we are mere observers, in which we cannot intervene: the heavens, for example. Since Newton, we have learned to apply the same causal laws to events in this realm; but in earlier times it was usually assumed that a quite different system of laws must operate there. It *could* have turned out that this was right; and then it could also have turned out that the system of laws we needed to explain events involving the celestial bodies required a notion of causality associated with the temporal direction from later to earlier.

When, however, we consider ourselves as agents, and consider causal laws governing events in which we can intervene, the notion of backwards causality seems to generate absurdities. If an event *C* is considered as the cause of a preceding event *D*, then it would be open to us to bring about *C* in order that the event *D* should have occurred. But the conception of doing something in order that something else should have happened appears to be intrinsically absurd: it apparently follows that backwards causation must also be absurd in any realm in which we can operate as agents.

We can affect the future by our actions: so why can we not by our actions affect the past? The answer that springs to mind is this: you cannot *change* the past; if a thing has happened, it has happened, and you cannot make it not to have happened. This is, I am told,¹ the attitude of orthodox Jewish theologians to retrospective prayer. It is blasphemous to pray that something should *have* happened, for, although there are no limits to God's power, He cannot do what is logically impossible; it is logically impossible to alter the past, so to utter a retrospective prayer is to mock God by asking Him to perform a logical impossibility. Now I think it is helpful to think about this example, because it is the only instance of behaviour, on the part of ordinary people whose mental processes we can understand, designed to affect the past and coming quite naturally to us. If one does not think of this case, the idea of doing something in order that something else should previously have happened may seem sheer raving insanity. But suppose I hear on the radio that a ship has gone down in the Atlantic two hours previously, and that there were a few survivors: my son was on that ship, and I at once utter a prayer that he should have been among the survivors, that he should not have

¹ By Professor G. Kreisel.

drowned; this is the most natural thing in the world. Still, there are things which it is very natural to say which make no sense; there are actions which can naturally be performed with intentions which *could* not be fulfilled. Are the Jewish theologians right in stigmatizing my prayer as blasphemous?

They characterize my prayer as a request that, if my son has drowned, God should make him not have drowned. But why should they view it as asking anything more self-contradictory than a prayer for the future? If, before the ship set sail, I had prayed that my son should make a safe crossing, I should not have been praying that, if my son was going to drown, God should have made him not be going to drown. Here we stumble on a well-known awkwardness of language. There is a use of the future tense to express present tendencies: English newspapers sometimes print announcements of the form 'The marriage that was arranged between X and Y will not now take place'. If someone did not understand the use of the future tense to express present tendencies, he might be amazed by this 'now'; he might say, 'Of course it will not take place *now*: either it is taking place *now*, or it will take place *later*'. The presence of the 'now' indicates a use of the future tense according to which, if anyone had said earlier, 'They are going to get married,' he would have been right, even though their marriage never subsequently occurred. If, on the other hand, someone had offered a bet which he expressed by saying, 'I bet they will not be married on that date,' this 'will' would normally be understood as expressing the *genuine* future tense, the future tense so used that what happens on the future date is the decisive test for truth or falsity, irrespective of how things looked at the time of making the bet, or at any intervening time. The future tense that I was using, and that will be used throughout this paper, is intended to be understood as this genuine future tense.

With this explanation, I will repeat: when, before the ship sails, I pray that my son will make the crossing safely, I am not praying that God should perform the logically impossible feat of making what will happen not happen (that is, not be-going-to happen); I am simply praying that it will not happen. To put it another way: I am not asking God that He should now make what is going to happen not be going to happen; I am asking that He *will* at a future time make something not to happen at that time. And similarly with my retrospective prayer. Assuming that I am not asking for a miracle—asking that, if my son has died, he should now be brought to life again—I do not have to be asking for a logical impossibility. I am not asking God that, even if my son has drowned, He should *now* make him not to have drowned; I am asking that, at the time of the disaster, He should then have made my son not to drown at that time. The former interpretation would indeed be required if the list of survivors had been read out over the radio, my son's name had not

been on it, and I had not envisaged the possibility of a mistake on the part of the news service: but in my ignorance of whether he was drowned or not, my prayer will bear another interpretation.

But this still involves my trying to affect the past. On this second interpretation, I am trying by my prayer *now* to bring it about that God made something not to happen: and is not this absurd? In this particular case, I can provide a rationale for my action—that is why I picked this example—but the question can be raised whether it is not a bad example, on the ground that it is the only kind for which a rationale *could* be given. The rationale is this. When I pray for the future, my prayer makes sense because I know that, at the time about which I am praying, God will remember my prayer, and may then grant it. But God knows everything, both what has happened and what is going to happen. So my retrospective prayer makes sense, too, because, at the time about which I am praying, God knew that I was going to make this prayer, and may then have granted it. So it seems relevant to ask whether foreknowledge of this kind can meaningfully be attributed only to God, in which case the example will be of a quite special kind, from which it would be illegitimate to generalize, or whether it could be attributed to human beings, in which case our example will not be of purely theological interest.

I have heard three opinions expressed on this point. The first, held by Russell and Ayer, is that foreknowledge is simply the mirror image of memory, to be explained in just the same words as memory save that 'future' replaces 'past', and so forth, and as such is conceptually unproblematic: we do not have the faculty but we perfectly well might. The second is a view held by a school of Dominican theologians. It is that God's knowledge of the future should be compared rather to a man's knowledge of what is going to happen, when this lies in his intention to make it happen. For example, God knows that I am going to pray that my son may not have drowned because He is going to make me pray so. This leads to the theologically and philosophically disagreeable conclusion that everything that happens is directly effected by God, and that human freedom is therefore confined to wholly interior movements of the will. This is the view adopted by Wittgenstein in the *Tractatus*, and there expressed by the statement 'The world is independent of my will.' On this view, God's foreknowledge is knowledge of a type that human beings do have; it would, however, be difficult to construct a non-theological example of an action intelligibly designed to affect the past by exploiting this alleged parallelism. The third view is one of which it is difficult to make a clear sense. It is that foreknowledge is something that can be meaningfully ascribed only to God (or perhaps also to those He directly inspires, the prophets; but again, perhaps these would be regarded not as themselves possessing this knowledge, but only as the instruments of

its expression). The ground for saying this is that the future is not something of which we could, but merely do not happen to, have knowledge; it is not, as it were, *there* to be known. Statements about the future are, indeed, either-true-or-false; but they do not yet have a particular one of these two truth-values. They have present truth-or-falsity, but they do not have present truth or present falsity, and so they *cannot* be known: there is not really anything to be known. The non-theological part of this view seems to me to rest on a philosophical confusion; the theological part I cannot interpret, since it appears to involve ascribing to God the performance of a logical impossibility.

We saw that retrospective prayer does not involve asking God to perform the logically impossible feat of changing the past, any more than prayer for the future involves asking Him to change the future in the sense in which that is logically impossible. We saw also that we could provide a rationale for retrospective prayer, a rationale which depended on a belief in God's foreknowledge. This led us to ask if foreknowledge was something which a man could have. If so, then a similar rationale could be provided for actions designed to affect the past, when they consisted in my doing something in order that someone should have known that I was going to do it, and should have been influenced by this knowledge. This enquiry, however, I shall not pursue any further. I turn instead to more general considerations: to consider other arguments designed to show an intrinsic absurdity in the procedure of attempting to affect the past—of doing something in order that something else should have happened. In the present connection I remark only that, if there is an intrinsic absurdity in *every* procedure of this kind, then it follows indirectly that there is also an absurdity in the conception of foreknowledge, human or divine.

Suppose someone were to say to me, 'Either your son has drowned or he has not. If he has drowned, then certainly your prayer will not (cannot) be answered. If he has not drowned, your prayer is superfluous. So in either case your prayer is pointless: it cannot make any *difference* to whether he has drowned or not.' This argument may well appear quite persuasive, until we observe that it is the exact analogue of the standard argument for fatalism. I here characterize fatalism as the view that there is an intrinsic absurdity in doing something in order that something else should subsequently happen; that any such action—that is, any action done with a further purpose—is necessarily pointless. The standard form of the fatalist argument was very popular in London during the bombing. The siren sounds, and I set off for the air-raid shelter in order to avoid being killed by a bomb. The fatalist argues, 'Either you are going to be killed by a bomb or you are not going to be. If you are, then any precautions you take will be ineffective. If you are

not, all precautions you take are superfluous. Therefore it is pointless to take precautions.' This belief was extended even to particular bombs. If a bomb was going to kill me, then it 'had my number on it', and there was no point in my attempting to take precautions against being killed by *that* bomb; if it did not have my number on it, then of course precautions were pointless too. I shall take it for granted that no one wants to accept this argument as cogent. But the argument is formally quite parallel to the argument supposed to show that it is pointless to attempt to affect the past; only the tenses are different. Someone may say, 'But it is just the difference in tense that makes the difference between the two arguments. Your son has either *already* been drowned or else *already* been saved; whereas you haven't *yet* been killed in the raid, and you haven't *yet* come through it.' But this is just to reiterate that the one argument is about the past and the other about the future: we want to know what, if anything, there is *in* this fact which makes the one valid, the other invalid. The best way of asking this question is to ask, 'What refutation is there of the fatalist argument, to which a quite parallel refutation of the argument to show that we cannot affect the past could not be constructed?'

Let us consider the fatalist argument in detail. It opens with a tautology, 'Either you are going to be killed in this raid or you are not'. As is well known, some philosophers have attempted to escape the fatalist conclusion by faulting the argument at this first step, by denying that two-valued logic applies to statements about future contingents. Although this matter is worth investigating in detail, I have no time to go into it here, so I will put the main point very briefly. Those who deny that statements about future contingents need be either true or false are under the necessity to explain the meaning of those statements in some way; they usually attempt to do so by saying something like this: such a statement is not true or false now, but *becomes* true or false at the time to which it refers. But if this is said, then the fatalist argument can be reconstructed by replacing the opening tautology by the assertion 'Either the statement "You will be killed in this raid" is going to become true, or it is going to become false'. The only way in which it can be consistently maintained not only that the law of excluded middle does not hold for statements about the future, but that there is no other logically necessary statement which will serve the same purpose of getting the fatalist argument off the ground, is to deny that there is, or could be, what I called a 'genuine' future tense at all: to maintain that the only intelligible use of the future tense is to express present tendencies. I think that most people would be prepared to reject this as unacceptable, and here, for lack of space, I shall simply assume that it is. (In fact, it is not quite easy to refute someone who consistently adopts this position; of course, it is always much easier to

make out that something is not meaningful than to make out that it is.) Thus, without more ado, I shall set aside the suggestion that the flaw in the fatalist argument lies in the very first step.

The next two steps stand or fall together. They are: 'If you are going to be killed in this raid, you will be killed whatever precautions you take' and 'If you are not going to be killed in this raid, you will not be killed whatever precautions you neglect'. These are both of the form 'If *p*, then if *q* then *p*'; for example, 'If you *are* going to be killed, then you will be killed even if you take precautions'. They are clearly correct on many interpretations of 'if'; and I do not propose to waste time by enquiring whether they are correct on 'the' interpretation of 'if' proper to well-instructed users of the English language. The next two lines are as follows: 'Hence, if you are going to be killed in the raid, any precautions you take will be ineffective' and 'Hence, if you are not going to be killed in the raid, any precautions you take will have been superfluous'. The first of these is indisputable. The second gives an appearance of sophistry. The fatalist argues from 'If you are not going to be killed, then you won't be killed even if you have taken no precautions' to 'If you are not going to be killed, then any precautions you take will have been superfluous'; that is, granted the truth of the statement 'You will not be killed even if you take no precautions', you will have no motive to take precautions; or, to put it another way, if you would not be killed even if you took no precautions, then any precautions you take cannot be considered as being effective in bringing about your survival—that is, as effecting it. This employs a well-known principle. St Thomas, for instance, says it is a condition of ignorance to be an excuse for having done wrong that, if the person had not suffered from the ignorance, he would not have committed the wrongful act in question. But we want to object that it may be just the precautions that I am going to take which save me from being killed; so it cannot follow from the mere fact that I am not going to be killed that I should not have been going to be killed even if I had not been going to take precautions. Here it really does seem to be a matter of the way in which 'if' is understood; but, as I have said, I do not wish to call into question the legitimacy of a use of 'if' according to which '(Even) if you do not take precautions, you will not be killed' follows from 'You will not be killed'. It is, however, clear that, on any use of 'if' on which this inference is valid, it is possible that both of the statements 'If you do not take precautions, you will be killed' and 'If you do not take precautions, you will not be killed' should be true. It indeed follows from the truth of these two statements together that their common antecedent is false; that is, that I am in fact going to take precautions. (It may be held that on a, or even the, use of 'if' in English, these two statements cannot both be true; or again, it may be held

that they can both be true only when a stronger consequence follows, namely, that not only am I as a matter of fact going to take precautions, but that I could not fail to take them, that it was not in my power to refrain from taking them. But, as I have said, it is not my purpose here to enquire whether there are such uses of 'if' or whether, if so, they are important or typical uses.) Now let us say that it is correct to say of certain precautions that they are capable of being effective in preventing my death in the raid if the two conditional statements are true that, if I take them, I shall not be killed in the raid, and that, if I do not take them, I shall be killed in the raid. Then, since, as we have seen, the truth of these two statements is quite compatible with the truth of the statement that, if I do not take precautions, I shall not be killed, the truth of this latter statement cannot be a ground for saying that my taking precautions will not be effective in preventing my death.

Thus, briefly, my method of rebutting the fatalist is to allow him to infer from 'You will not be killed' to 'If you do not take precautions, you will not be killed'; but to point out that, on any sense of 'if' on which this inference is valid, it is impermissible to pass from 'If you do not take precautions, you will not be killed' to 'Your taking precautions will not be effective in preventing your death'. For this to be permissible, the truth of 'If you do not take precautions, you will not be killed' would have to be incompatible with that of 'If you do not take precautions, you will be killed'; but, on the sense of 'if' on which the first step was justified, these would not be incompatible. I prefer to put the matter this way than to make out that there is a sense of 'if' on which these two are indeed incompatible, but on which the first step is unjustified, because it is notoriously difficult to elucidate such a sense of 'if'.

Having arrived at a formulation of the fallacy of the fatalist argument, let us now consider whether the parallel argument to demonstrate the absurdity of attempting to bring about the past is fallacious in the same way. I will abandon the theological example in favour of a magical one. Suppose we come across a tribe who have the following custom. Every second year the young men of the tribe are sent, as part of their initiation ritual, on a lion hunt: they have to prove their manhood. They travel for two days, hunt lions for two days, and spend two days on the return journey; observers go with them, and report to the chief upon their return whether the young men acquitted themselves with bravery or not. The people of the tribe believe that various ceremonies, carried out by the chief, influence the weather, the crops, and so forth. I do not want these ceremonies to be thought of as religious rites, intended to dispose the gods favourably towards them, but simply as performed on the basis of a wholly mistaken system of causal beliefs. While the young men are away from the village the chief performs ceremonies—

dances, let us say—intended to cause the young men to act bravely. We notice that he continues to perform these dances for the whole six days that the party is away, that is to say, for two days during which the events that the dancing is supposed to influence have already taken place. Now there is generally thought to be a *special* absurdity in the idea of affecting the past, much greater than the absurdity of believing that the performance of a dance can influence the behaviour of a man two days' journey away; so we ought to be able to persuade the chief of the absurdity of his continuing to dance after the first four days without questioning his general system of causal beliefs. How are we going to do it?

Since the absurdity in question is alleged to be a *logical* absurdity, it must be capable of being seen to be absurd however things turn out; so I am entitled to suppose that things go as badly for us, who are trying to persuade the chief of this absurdity, as they can do; we ought still to be able to persuade him. We first point out to him that he would not think of continuing to perform the dances after the hunting party has returned; he agrees to that, but replies that that is because at that time he *knows* whether the young men have been brave or not, so there is no longer any point in trying to bring it about that they have been. It is irrelevant, he says, that during the last two days of the dancing they have already either been brave or cowardly: there is still a point in his trying to make them have been brave, because he does not yet know which they have been. We then say that it can be only the first four days of the dancing which could possibly affect the young men's performance; but he replies that experience is against that. There was for several years a chief who thought as we did, and danced for the first four days only; the results were disastrous. On two other occasions, he himself fell ill after four days of dancing and was unable to continue, and again, when the hunting party returned, it proved that the young men had behaved ignobly.

The brief digression into fatalism was occasioned by our noticing that the standard argument against attempting to affect the past was a precise analogue of the standard fatalist argument against attempting to affect the future. Having diagnosed the fallacy in the fatalist argument, my announced intention was to discover whether there was not a similar fallacy in the standard argument against affecting the past. And it indeed appears to me that there is. We say to the chief, 'Why go on dancing now? Either the young men have already been brave, or they have already been cowardly. If they have been brave, then they have been brave whether you dance or not. If they have been cowardly, then they have been cowardly whether you dance or not. If they have been brave, then your dancing now will not be effective in making them have been brave, since they have been brave even if you do not dance. And if they have not been brave, then your dancing will certainly not be effective.

Thus your continuing to dance will in the one case be superfluous, and in the other fruitless: in neither case is there any point in your continuing to dance.' The chief can reply in exactly the way in which we replied to the fatalist. He can say, 'If they have been brave, then indeed there is a sense in which it will be true to say that, even if I do not dance, they will have been brave; but this is not incompatible with its also being true to say that, if I do not dance, they will not have been brave. Now what saying that my continuing to dance is effective in causing them to have been brave amounts to is that it is true both that, if I go on dancing, they have been brave, and that, if I do not dance, they have not been brave. I have excellent empirical grounds for believing both these two statements to be true; and neither is incompatible with the truth of the statement that, if I do not dance, they have been brave, although, indeed, I have no reason for believing *that* statement. Hence, you have not shown that, from the mere hypothesis that they have been brave, it follows that the dancing I am going to do will not be effective in making them have been brave; on the contrary, it may well be that, although they have been brave, they have been brave just *because* I am going to go on dancing; that, if I were not going to go on dancing, they would not have been brave.' This reply sounds sophistical; but it cannot be sophistical if our answer to the fatalist was correct, because it is the exact analogue of that answer.

We now try the following argument: 'Your *knowledge* of whether the young men have been brave or not may affect whether you *think* there is any point in performing the dances; but it cannot really make any difference to the *effect* the dances have on what has happened. If the dances are capable of bringing it about that the young men have acted bravely, then they ought to be able to do that even after you have learned that the young men have *not* acted bravely. But that is absurd, for that would mean that the dances can change the past. But if the dances cannot have any effect after you have learned whether the young men have been brave or not, they cannot have any effect before, either; for the mere state of your knowledge cannot make any difference to their efficacy.' Now since the causal beliefs of this tribe are so different from our own, I could imagine that the chief might simply deny this: he might say that what had an effect on the young men's behaviour was not merely the performance of the dances by the chief as such, but rather their performance by the chief when in a state of ignorance as to the outcome of the hunt. And if he says this, I think there is really no way of dissuading him, short of attacking his whole system of causal beliefs. But I will not allow him to say this, because it would make his causal beliefs so different in kind from ours that there would be no moral to draw for our own case. Before going on to consider his reaction to this argument, however, let us first pause to review the situation.

Suppose, then, that he agrees to our suggestion: agrees, that is, that it is his dancing as such that he wants to consider as bringing about the young men's bravery, and not his dancing in ignorance of whether they were brave. If this is his belief, then we may reasonably challenge him to try dancing on some occasion when the hunting party has returned and the observers have reported that the young men have *not* been brave. Here at last we appear to have hit on something which has no parallel in the case of affecting the future. If someone believes that a certain kind of action is effective in bringing about a subsequent event, I may challenge him to try it out in all possible circumstances: but I cannot demand that he try it out on some occasion when the event is *not* going to take place, since he cannot identify any such occasion independently of his intention to perform the action. Our knowledge of the future is of two kinds: prediction based on causal laws and knowledge in intention. If I think I can predict the non-occurrence of an event, then I cannot consistently also believe that I can do anything to bring it about; that is, I cannot have good grounds for believing, of any action, both that it is in my power to do it, and that it is a condition of the event's occurring. On the other hand, I cannot be asked to perform the action on some occasion when I believe that the event will not take place, when this knowledge lies in my intention to prevent it taking place; for as soon as I accede to the request, I thereby abandon my intention. It would, indeed, be different if we had foreknowledge: someone who thought, like Russell and Ayer, that it is a merely contingent fact that we have memory but not foreknowledge would conclude that the difference I have pointed to does not reveal a genuine asymmetry between past and future, but merely reflects this contingent fact.

If the chief accepts the challenge, and dances when he knows that the young men have not been brave, it seems that he must concede that his dancing does not *ensure* their bravery. There is one other possibility favourable to us. Suppose that he accepts the challenge, but when he comes to try to dance, he unaccountably cannot do so: his limbs simply will not respond. Then we may say, 'It is not your dancing (after the event) which causes them to have been brave, but rather their bravery which makes possible your dancing: your dancing is not, as you thought, an action which it is in your power to do or not to do as you choose. So you ought not to say that you dance in the last two days in order to make them have been brave, but that you try to see whether you can dance, in order to find out whether they have been brave.'

It may seem that this is conclusive; for are not these the only two possibilities? Either he does dance, in which case the dancing is proved not to be a sufficient condition of the previous bravery; or he does not, in which case

the bravery must be thought a causal condition of the dancing rather than vice versa. But in fact the situation is not quite so simple.

For one thing, it is not justifiable to demand that the chief should either consider his dancing to be a sufficient condition of the young men's bravery, or regard it as wholly unconnected. It is enough, in order to provide him with a motive for performing the dances, that he should have grounds to believe that there is a significant positive correlation between his dancing and previous brave actions on the part of the young men; so the occurrence of a certain proportion of occasions on which the dancing is performed, although the young men were not brave, is not a sufficient basis to condemn him as irrational if he continues to dance during the last two days. Secondly, while his being afflicted with an otherwise totally inexplicable inability to dance may strongly suggest that the cowardice of the young men renders him unable to dance, and that therefore dancing is not an action which it is in his power to perform as he chooses, any failure to dance that is explicable without reference to the outcome of the hunt has much less tendency to suggest this. Let us suppose that we issue our challenge, and he accepts it. On the first occasion when the observers return and report cowardly behaviour on the part of the young men, he performs his dance. This weakens his belief in the efficacy of the dancing, but does not disturb him unduly; there have been occasions before when the dancing has not worked, and he simply classes this as one of them. On the second occasion when the experiment can be tried, he agrees to attempt it, but, a few hours before the experiment is due to be carried out, he learns that a neighbouring tribe is marching to attack his, so the experiment has to be abandoned; on the third occasion, he is bitten by a snake, and so is incapacitated for dancing. Someone might wish to say, 'The cowardice of the young men caused those events to happen and so prevent the chief from dancing', but such a description is far from mandatory: the chief may simply say that these events were accidental, and in no way *brought about* by the cowardice of the young men. It is true that if the chief is willing to attempt the experiment a large number of times, and events of this kind repeatedly occur, it will no longer appear reasonable to dismiss them as a series of coincidences. If accidents which prevent his dancing occur on occasions when the young men are known to have been cowardly with much greater frequency than, say, in a control group of dancing attempts, when the young men are known to have been brave, or when it is not known how they behaved, then this frequency becomes something that must itself be explained, even though each particular such event already has its explanation.

Suppose now, however, that the following occurs. We ask the chief to perform the dances on some occasion when the hunting party has returned

and the observers have reported that the young men have not acquitted themselves with bravery. He does so, and we claim another weakening of his belief that the dancing is correlated with preceding bravery. But later it turns out that, for some reason or other, the observers were lying (say they had been bribed by someone): so after all this is not a counter-example to the law. So we have a third possible outcome. The situation now is this. We challenge the chief to perform the dances whenever he knows that the young men have not been brave, and he accepts the challenge. There are three kinds of outcome: (i) he simply performs the dances; (ii) he is prevented from performing the dances by some occurrence which has a quite natural explanation totally independent of the behaviour of the young men; and (iii) he performs the dances, but subsequently discovers that this was not really an occasion on which the young men had not been brave. We may imagine that he carries out the experiment repeatedly, and that the outcome always falls into one of these three classes; and that outcomes of class (i) are sufficiently infrequent not to destroy his belief that there is a significant correlation between the dancing and the young men's bravery, and outcomes of class (ii) sufficiently infrequent not to make him say that the young men's cowardice renders him incapable of performing the dances. Thus our experiment has failed.

On the other hand, it has not left everything as before. I have exploited the fact that it is frequently possible to discover that one had been mistaken in some belief about the past. I will not here raise the question whether it is *always* possible to discover this, or whether there are beliefs about the past about which we can be *certain* in the sense that nothing could happen to show the belief to have been mistaken. Now before we challenged the chief to perform this series of experiments, his situation was as follows. He was prepared to perform the dancing in order to bring it about that the young men had been brave, but only when he had no information about whether they had been brave or not. The rationale of his doing so was simply this: experience shows that there is a positive correlation between the dancing and the young men's bravery; hence the fact that the dances are being performed makes it more probable that the young men have been brave. But the dancing is something that is in my power to do if I choose: experience does not lead me to recognize it as a possibility that I should try to perform the dances and fail. Hence it is in my power to do something, the doing of which will make it more probable that the young men have been brave: I have therefore every motive to do it. Once he had information, provided by the observers, about the behaviour of the young men, then, under the old dispensation, his attitude changed: he no longer had a motive to perform the dances. We do not have to assume that he was unaware of the possibility that the observers were lying

or had made a mistake. It may just have been that he reckoned the probability that they were telling the truth as so high that the performance of the dances after they had made their report would make no significant difference to the probability that the young men had been brave. If they reported the young men as having been brave, there was so little chance of their being wrong that it was not worth while to attempt to diminish this chance by performing the dances; if they reported that the young men had been cowardly, then even the performance of the dances would still leave it overwhelmingly probable that they *had* been cowardly. That is to say, until the series of experiments was performed, the chief was prepared to discount completely the probability conferred by his dancing on the proposition that the young men had been brave in the face of a source of information as to the truth of this proposition of the kind we ordinarily rely upon in deciding the truth or falsity of statements about the past. And the reason for this attitude is very clear: for the proposition that there was a positive correlation between the dancing and the previous bravery of the young men could have been established in the first place only by relying on our ordinary sources of information as to whether the young men had been brave or not.

But if we are to suppose that the series of experiments works out in such a way as not to force the chief to abandon his belief both that there is such a positive correlation and that the dancing is something which it is in his power to do when he chooses, we must suppose that it fairly frequently happens that the observers are subsequently proved to have been making false statements. And I think it is clear that in the process the attitude of the chief to the relative degree of probability conferred on the statement that the young men have been brave by (i) the reports of the observers and (ii) his performance of the dances will alter. Since it so frequently happens that, when he performs the dances *after* having received an adverse report from the observers, the observers prove to have been misreporting, he will cease to think it pointless to perform the dances after having received such an adverse report: he will thus cease to think that he can decide whether to trust the reports of the observers independently of whether he is going to perform the dances or not. In fact, it seems likely that he will come to think of the performance of the dances as itself a ground for distrusting, or even for denying outright, the adverse reports of the observers, even in the absence of any *other* reason (such as the discovery of their having been bribed, or the reports of some other witness) for believing them not to be telling the truth.

The chief began with two beliefs: (i) that there was a positive correlation between his dancing and the previous brave behaviour of the young men; and (ii) that the dancing was something in his power to do as he chose. We are tempted to think of these two beliefs as incompatible, and I described people

attempting to devise a series of experiments to convince the chief of this. I tried to show, however, that these experiments could turn out in such a way as to allow the chief to maintain both beliefs. But in the process a third belief, which we naturally take for granted, has had to be abandoned in order to hang on to the first two: the belief, namely, that it is possible for me to find out what has happened (whether the young men have been brave or not) independently of my intentions. The chief no longer thinks that there is any evidence as to whether the young men had been brave or not, the strength of which is unaffected by whether he intends subsequently to perform the dances. And now it appears that there really is a form of incompatibility among these *three* beliefs, in the sense that it is always possible to carry out a series of actions which will necessarily lead to the abandonment of at least one of them. Here there is an exact parallel with the case of affecting the future. We *never* combine the beliefs (i) that an action *A* is positively correlated with the subsequent occurrence of an event *B*; (ii) that the action *A* is in my power to perform or not as I choose; and (iii) that I can know whether *B* is going to take place or not independently of my intention to perform or not to perform the action *A*. The difference between past and future lies in this: that we think that, of any past event, it is in principle possible for me to know whether or not it took place independently of my present intentions; whereas, for many types of future event, we should admit that we are never going to be in a position to have such knowledge independently of our intentions. (If we had foreknowledge, this might be different.) If we insist on hanging on to this belief, for all types of past event, then we cannot combine the two beliefs that are required to make sense of doing something in order that some event should have previously taken place; but I do not know any reason why, if things were to turn out differently from the way they do now, we *could* not reasonably abandon the first of these beliefs rather than either of the other two.

My conclusion therefore is this. If anyone were to claim, of some type of action *A*, (i) that experience gave grounds for holding the performance of *A* as increasing the probability of the previous occurrence of a type of event *E*; and (ii) that experience gave no grounds for regarding *A* as an action which it was ever not in his power to perform—that is, for entertaining the possibility of his trying to perform it and failing—then we could either force him to abandon one or other of these beliefs, or else to abandon the belief (iii) that it was ever possible for him to have knowledge, independent of his intention to perform *A* or not, of whether an event *E* had occurred. Now doubtless most normal human beings would rather abandon either (i) or (ii) than (iii), because we have the prejudice that (iii) must hold good for every type of event: but if someone were, in a particular case, more ready to give

up (iii) than (i) or (ii), I cannot see any argument we could use to dissuade him. And so long as he was not dissuaded, he could sensibly speak of performing *A* in order that *E* should have occurred. Of course, he could adopt an intermediate position. It is not really necessary, for him to be able to speak of doing *A* in order that *E* should have occurred, that he deny all possibility of his trying and failing to perform *A*. All that is necessary is that he should not regard his being informed, by ordinary means, of the non-occurrence of *E* as making it more probable that if he tries to perform *A*, he will fail: for, once he does so regard it, we can claim that he should regard the occurrence of *E* as making possible the performance of *A*, in which case his trying to perform *A* is not a case of trying to bring it about that *E* has happened, but of finding out whether *E* has happened. (Much will here depend on whether there is an ordinary causal explanation for the occurrence of *E* or not.) Now he need not really deny that learning, in the ordinary way, that *E* has not occurred makes it at all more probable that, if he tries to perform *A*, he will fail. He may concede that it makes it to some extent more probable, while at the same time maintaining that, even when he has grounds for thinking that *E* has not occurred, his intention to perform *A* still makes it more probable than it would otherwise be that *E* has in fact occurred. The attitude of such a man seems paradoxical and unnatural to us, but I cannot see any rational considerations which would force him out of this position. At least, if there are any, it would be interesting to know what they are: I think that none of the considerations I have mentioned in this paper could serve this purpose.

My theological example thus proves to have been a bad—that is, untypical—example in a way we did not suspect at the time, for it will never lead to a discounting of our ordinary methods of finding out about the past. I may pray that the announcer has made a mistake in not including my son's name on the list of survivors; but once I am convinced that no mistake has been made, I will not go on praying for him to have survived. I should regard this kind of prayer as something to which it was possible to have recourse only when an ordinary doubt about what had happened could be entertained. But just because this example is untypical in this way, it involves no tampering with our ordinary conceptual apparatus at all: this is why it is such a natural thing to do. On my view, then, orthodox Jewish theology is mistaken on this point.

I do not know whether it could be held that part of what people have meant when they have said 'You cannot change the past' is that, for every type of event, it is in principle possible to know whether or not it has happened, independently of one's own intentions. If so, this is not the mere tautology it appears to be, but it does indeed single out what it is that makes us think it impossible to bring about the past.