


11

The Powers of Rational Beings:
Freedom of the Will

WE NOW TURN TO another mystery, a mystery about the *powers* of rational beings; that is, a mystery about what human beings are able to do. This mystery is the mystery of free will and determinism. The best way to get an intuitive grip on the problem of free will and determinism is to think of time as a "garden of forking paths." That is, to think of the alternatives that one considers when one is deciding what to do as being parts of various "alternative futures" and to think of these alternative futures diagrammatically, in the way suggested by a path or a river or a road that literally forks:




If Jane is trying to decide whether to tell all or to continue her life of deception, she is in a situation strongly analogous to that of someone who is hesitating between forks in a road. That is why this sort of diagram is so suggestive. Let us apply this idea to the problem of free will and determinism.


To say that one has free will is to say that when one decides among forks in the road of time (or, more prosaically, when one decides what to do), one is at least sometimes able to take more than one of the forks. Thus, Jane, who is deciding between a fork that leads to telling all and a fork that leads to a life of continued deception, has free will (on this particular occasion) if she is able to tell all and is also able to continue living a life of deception. One has free will if sometimes more than one of the forks in the road of time is "open" to one. One lacks free will if on every occasion on which one must make a decision only one of the forks before one—of course it will be the fork one in fact takes—is open to one. If John is locked in a room and doesn't know that he is locked in, and if he is in the process of deliberating about whether to leave, one of the alternative futures he is contemplating—leaving—is, in point of fact, not open to him, and he thus lacks free will in the matter of staying or leaving.¹

It is a common opinion that free will is required by morality. Let us examine this common opinion from the perspective that is provided by looking at time as a garden of forking paths. While it is obviously false—for about six independent reasons—that the whole of morality consists in making judgments of the form 'You should not have done X', we can at least illustrate certain important features of the relation between free will and morality by examining the relation between the concept of free will and the content of such judgments. The judgment that you shouldn't have done X implies that you should have done something else instead; that you should have done something else instead implies that there was something else for you to do; that there was something else for you to do implies that you *could* have done something else; that you could have done something else implies that you have free will. To make a moral judgment about one of your acts is to evaluate your taking one of the forks in the road of time, to characterize it as better or worse than various of the other forks that were open to you. (Note that if you have made a choice by taking one of the forks in what is literally a road, no one could blame you for taking the fork you did if all of the other forks were blocked.) A moral evaluation of what someone has done requires two or more alternative possibilities of action for that person just as surely as a contest requires two or more contestants.

Let us now see what help the conception of time as a garden of forking paths gives us in understanding what is meant by determinism. Determinism is the thesis that it is true at every moment that the way things then are determines a unique future, that only *one* of the alternative futures that may exist relative to a given moment is a physically possible continuation of the state of things at that moment. Or, if you like, we may say that determinism is the thesis that only one continuation of the state of things at a given moment is consistent with the laws of nature. (For it is the laws of nature that determine what is physically possible. It is, for example, now physically possible for you to be in Chicago at noon tomorrow if and only if your being in Chicago at noon tomorrow is consistent with both the present state of things and the laws of nature.) Thus, according to determinism, although it may often seem to us that we confront a sheaf of possible futures (like this)



what we really confront is something like this



This figure is almost shaped like a road that splits into four roads, but not quite: three of the four "branches" that lead away from the "fork" are not connected with the original road, although they come very close to it. (Thus they are not really branches in the road, and the place at which they almost touch the road is not really a fork.) If we were to view this figure from a distance—across the room, say—it would seem to us to have the shape of a road that forks. We have to look at it closely to see that what appeared from a distance to be three "branches" are not connected with the long line or with one another. In the figure, the point at which the three unconnected lines *almost* touch the long line represents the present. The unconnected lines represent futures that are not physically possible continuations of the present, and the part of the long line to the right of the "present" represents a future that is a physically possible continuation of the present. The gaps between the long line and the unconnected lines represent causal discontinuities, violations of the laws of nature—in a word, miracles. The reason these futures are not physically possible continuations of the present is that "getting into" any of them from the present would require a miracle. The fact that the part of the long line that lies to the right of the "present" actually proceeds from that point represents the fact that this line-segment corresponds to a physically possible future.

This figure, then, represents four futures, three of which are physically impossible and exactly one of which is physically possible. If these four futures are the only futures that "follow" the present, then this figure represents the way in which each moment of time must be if the universe is deterministic: each moment must be followed by exactly one physically possible future.

The earlier diagram, however, represents an indeterministic situation. The road really does fork. The present is followed by four possible futures. Any one of them could, consistently with the laws of nature, evolve out of the present. Any one of them could, consistently with the laws of nature, turn out to be the actual future. Therefore, it is only if the universe is indeterministic that time *really is* a "garden of forking paths." But even in a deterministic universe, time could *look like* a garden of forking paths. Remember that our figure, when viewed from across the room, *looked* as if it had the shape of a road that forked. We cannot see all, or even very many, of the causes that operate in any situation. It could be, therefore, that the universe is deterministic, even though it looks to our limited vision as if there were sometimes more than one possible future. It may look to Jane as if she faces two possible futures, in one of which she tells all and in the other of which she continues her life of deception. But it may well be that the possibility of one or the other of these contemplated futures is mere appearance—an illusion, in fact. It may be that, in reality, causes already at work in her brain and central nervous system and immediate environment have already "ruled out" one or the other of these futures: it may be that one or the other of them is such that it could not come to pass unless a physically impossible event, a miracle, were to happen in her brain or central nervous system or environment.

Ask yourself this question. What would happen if some supernatural agency—God, say—were to "roll history back" to some point in the past and then "let things go forward again"? Suppose the agency were to cause things

to be once more just as they were at high noon, Greenwich time, on 11 March 1893 and were thereafter to let things go on of their own accord. Would history literally repeat itself? Would there be two world wars, each the same in every detail as the wars that occurred the "first time around"? Would a president of the United States called 'John F. Kennedy' be assassinated in Dallas on the date that on the new reckoning is called '22 November 1963'? Would you, or at least someone exactly like you, exist? If the answer to these questions is No, then determinism is false. Equivalently, if determinism is true, the answer to these questions is Yes. If determinism is true, then, if the universe were rolled back to a previous state by a miracle, and if there were no further miracles, the history of the world would repeat itself. And if the universe were rolled back to a previous state thousands of times, this exact duplication would happen every time. If there are no forks in the road of time—if all of the apparent forks are merely apparent, illusions due to our limited knowledge of the causes of things—then restoring the universe to some earlier condition is like moving a traveler on a road without forks back to an earlier point on that road. If there are no forks in the road, then, obviously enough, the traveler must traverse the same path a second time.

It has seemed obvious to most people who have not been exposed (perhaps 'subjected' would be a better word) to philosophy that free will and determinism are incompatible. It is almost impossible to get beginning students of philosophy to take seriously the idea that there could be such a thing as free will in a deterministic universe. Indeed, people who have not been exposed to philosophy usually understand the word 'determinism' (if they know the word at all) to stand for the thesis that there is no free will. And you might think that the incompatibility of free will and determinism deserves to seem obvious—because it is obvious. To say that we have free will is to say that more than one future is sometimes open to us. To affirm determinism is to say that every future that confronts us but one is physically impossible. And, surely, a physically impossible future can't be open to anyone, can it? If we know that a "Star Trek" sort of future is physically impossible (because, say, the "warp drives" and "transporter beams" that figure essentially in such futures are physically impossible), then we know that a "Star Trek" future is not open to us or to our descendants.

People who are convinced by this sort of reasoning are called *incompatibilists*: they hold that free will and determinism are incompatible. As I have hinted, however, many philosophers are *compatibilists*: they hold that free will and determinism are compatible. Compatibilism has an illustrious history among English-speaking philosophers, a history that embraces such figures as the seventeenth-century English philosopher Thomas Hobbes, the eighteenth-century Scottish philosopher David Hume, and the nineteenth-century English philosopher John Stuart Mill. And the majority of twentieth-century English-speaking philosophers have been compatibilists. (But compatibilism has not had many adherents on the continent of Europe. Kant, for example, called it a "wretched subterfuge.")

A modern compatibilist can be expected to reply to the line of reasoning I have just presented in some such way as follows:

Yes, a future, in order to be open to one, does need to be physically possible. It can't, for example, contain faster-than-light travel if faster-than-light travel is physically impossible. But we must distinguish between a future's being physically possible and its having a physically possible connection with the present. A future is physically possible if everything that happens in it is permitted by the laws of nature. A future has a physically possible connection with the present if it could be 'joined' to the present without any violation of the laws of nature. A physically possible future that does not have a physically possible connection with the present is one that, given the present state of things, would have to be 'inaugurated' by a miracle, an event that violated the laws of nature, but in which, thereafter, events proceeded in accordance with the laws. Determinism indeed says that of all the physically possible futures, one and only one has a physically possible connection with the present—one and only one could be joined to the present without a violation of the laws of nature. My position is that some futures that could not be joined to the present without a violation of the laws of nature are, nevertheless, open to us.

Two philosophical problems face the defenders of compatibilism. The easier is to provide a clear statement of *which* futures that do not have a physically possible connection with the present are "open" to us. The more difficult is to make it seem at least plausible that futures that are in this sense open to an agent really deserve to be so described.

An example of a solution to these problems may make the nature of the problems clearer. The solution I shall briefly describe would almost certainly be regarded by all present-day compatibilists as defective, although it has a respectable history. I choose it not to suggest that compatibilists can't do better but simply because it can be described in fairly simple terms.

According to this solution, a future is open to an agent, if, given that the agent chose that future (chose that path leading away from a fork in the road of time), it would come to pass. Thus it is open to me to stop writing this book and do a little dance because, if I so chose, that's what I'd do. But if Alice is locked in a prison cell, it is not open to her to leave: if she chose to leave, her choice would be ineffective because she would come up against a locked prison door. Now consider the future I said was open to me—to stop writing and do a little dance—and suppose that determinism is true. Although a choice on my part to behave in that remarkable fashion would (no doubt) be effective if it occurred, it is as a matter of fact *not* going to occur, and, therefore, given determinism, it is determined by the present state of things and the laws of nature that such a choice is not going to occur. It is in fact determined that *nothing* is going to occur that would have the consequence that I stop writing and do a little dance. Therefore, none of the futures in which I act in that bizarre way is a future that has a physically possible connection with the present: such a future could come to pass only if it were inaugurated by an event of a sort that is ruled out by the present state of things and the laws of nature.

And yet, as we have seen, many of these futures are "open" to me in the sense of 'open' that the compatibilist has proposed.

Is this a reasonable sense to give to this word? (We now take up the second problem that confronts the compatibilist.) This is a very large question. The core of the compatibilist's answer is an attempt to show that the reason we are interested in open or accessible futures is that we are interested in modifying the way people behave. One important way in which we modify behavior is by rewarding behavior that we like and punishing behavior that we dislike. We tell people that we will put them in jail if they steal and that they will get a tax break if they invest their money in such-and-such a way. But there is no point in trying to get people to act in a certain way if that way is not in some sense open to them. There is no point in telling Alfred that he will go to jail if he steals unless it is somehow open to him not to steal.

And what is the relevant sense of "open"? Just the one I have proposed, says the compatibilist. One modifies behavior by modifying the choices people make. That procedure is effective just insofar as choices are effective in producing behavior. If Alfred chooses not to steal (and remains constant in that choice), then he won't steal. But if Alfred chooses not to be subject to the force of gravity, he will nevertheless be subject to the force of gravity. Although it would no doubt be socially useful if there were some people who were not subject to the force of gravity, there is no point in threatening people with grave consequences if they do not break the bonds of gravity, for even if you managed to induce some people to choose not to be subject to the force of gravity, their choice would not be effective. Therefore (the compatibilist concludes), it is entirely appropriate to speak of a future as "open" if it is a future that would be brought about by a choice—even if it were a choice that was determined not to occur. And if Alfred protests when you punish him for not choosing a future that was in this sense open to him, on the ground that it was determined by events that occurred before his birth that he not make the choice that would have inaugurated that future—if he protests that only a *miracle* could have inaugurated such a future—you can tell him that his punishment will not be less effective in modifying his behavior (and the behavior of those who witness his punishment) on *that* account.

When things are put that way, compatibilism can look like nothing more than robust common sense. Why, then, do people have so much trouble believing it? Why does it arouse so much resistance? I think that the reason is that compatibilists can make their doctrine seem like robust common sense only by sweeping a mystery under the carpet and that, despite their best efforts, the bulge shows. People are aware that something is amiss with compatibilism even when they are unable to articulate their misgivings. I believe that it is possible to lift the carpet and display the hidden mystery. The notion of "not having a choice" has a certain logic to it. One of the principles of this logic is, or so it seems, embodied in the following thesis, which I shall refer to as the No Choice Principle:

Suppose that p and that no one has (or ever had) any choice about whether p . And suppose also that the following conditional (if-then) statement is

true and that no one has (or ever had) any choice about whether it is true: if p , then q . It follows from these two suppositions that q and that no one has (or ever had) any choice about whether q .

In this statement of the No Choice Principle, any declarative sentences can replace the symbols ' p ' and ' q '. (But the same sentence must replace ' p ' at each place it occurs, and the same goes for ' q '.) We might, for example, replace ' p ' with 'Plato died long before I was born' and ' q ' with 'I have never met Plato':

Suppose that Plato died long before I was born and that no one has (or ever had) any choice about whether Plato died long before I was born. And suppose also that the following conditional statement is true and that no one has (or ever had) any choice about whether it is true: if Plato died long before I was born, then I have never met Plato. It follows from these two suppositions that I have never met Plato and that no one has (or ever had) any choice about whether I have never met Plato.

The No Choice Principle seems undeniably correct. How could I have a choice about anything that is an inevitable consequence of something I have no choice about? And yet, as we shall see, the compatibilist must deny the No Choice Principle. To see why this is so, let us suppose that determinism is true and that the No Choice Principle is correct. Now let us consider some state of affairs that we should normally suppose someone had a choice about. Consider, say, the fact that I am writing this book. Most people—at least most people who knew I was writing a book—would assume that I had a choice about whether I was engaged in this project. They would assume that it was open to me to have undertaken some other project or no project at all. But we are supposing that determinism is true, and that means that ten million years ago (say) there was only one physically possible future, a future that included my being engaged in writing this book at the present date (since that is what I am in fact doing): given the way things were ten million years ago and given the laws of nature, it had to be true that I was now engaged in writing this book. But consider the two statements

- Things were thus-and-so ten million years ago.
- If things were thus-and-so ten million years ago, then I am working on this book now.

(Here 'thus-and-so' is a sort of gesture at a complete description or specification of the way things were ten million years ago.) Each of these statements is true. And it is obvious that no one has or ever had any choice about the truth of either. It is obvious that no one—no human being, certainly—has or ever had any choice about whether things were thus-and-so ten million years ago, since at that time the first human beings were still millions of years in the future.

And no one has any choice about whether the second statement, the if-then statement, is true because this statement is a consequence of the laws of nature, and no one—no human being, certainly—has any choice about what the

laws of nature are. If we imagine a possible world in which, as in the actual world, things were thus-and-so ten million years ago, and in which, unlike in the actual world, I decided to learn to sail instead of writing this book, we are imagining a world in which the laws of nature are different; for the *actual* laws dictate that if at some point in time things are thus-and-so, then, ten million years later I (or at any rate someone just like me) shall be writing and not sailing.

But if both of the above statements are true, then it follows, by the No Choice Principle, that neither I nor anyone else has or ever had any choice about whether I write this book. And, obviously, the content of the particular example—my writing a book—played no role in the derivation of this conclusion. It follows that, given the No Choice Principle, determinism implies that there is no free will. That is why the compatibilist must reject the No Choice Principle. This is the hidden mystery that, I contend, lies behind the façade of bluff common sense that compatibilism presents to the world: the compatibilist must reject the No Choice Principle, and the No Choice Principle seems to be true beyond all possibility of dispute. (Either that or the compatibilist must hold that one can have a choice about what went on in the world before there were any human beings or that one can have a choice about what the laws of nature are. But these alternatives look even more implausible than a rejection of the No Choice Principle.) If the No Choice Principle were false, that would be a great mystery indeed.

We must not forget, however, that mysteries really do exist. There are principles that are commonly held, and with good reason, to be false and whose falsity seems to be just as great a mystery as the falsity of the No Choice Principle would be. Consider, for example the principle that is usually called "the Galilean Law of the Addition of Velocities." This principle is a generalization of cases like the following. Suppose that an airplane is flying at a speed of 800 kilometers per hour relative to the ground; suppose that inside the aircraft a housefly is buzzing along at a speed of 30 kilometers per hour relative to the airplane in the direction of the airplane's travel; then the fly's speed relative to the ground is the sum of these two speeds: 830 kilometers per hour. According to the Special Theory of Relativity, an immensely useful and well-confirmed theory, the Galilean Law of the Addition of Velocities does not hold (although it comes very, very close to holding when it is applied to velocities of the magnitude that we usually consider in everyday life). And yet when one considers this principle in the abstract—in isolation from the considerations that guided Einstein in his development of Special Relativity—it seems to force itself upon the mind as true, to be true beyond all possibility of doubt. It seems, therefore, that the kind of "inner conviction" that sometimes moves one to say things like, "I can just see that that proposition *has* to be true" is not infallible.

Nevertheless, a mystery is a mystery. If compatibilism hides a mystery, should we therefore be incompatibilists? Unfortunately, incompatibilism also hides a mystery. Behold, I will show you a mystery.

If we are incompatibilists, we must reject either free will or determinism. What happens if we reject determinism? It is a bit easier now to reject determinism than it was in the nineteenth century, when it was commonly be-

lieved, and with reason, that determinism was underwritten by physics. But the quantum-mechanical world of current physics seems to be irreversibly indeterministic, and physics has therefore got out of the business of underwriting determinism. Nevertheless, the physical world is filled with objects and systems that seem to be deterministic "for all practical purposes"—digital computers, for example—and many philosophers and scientists believe that a human organism is deterministic for all practical purposes. But let us not debate this question. Let us suppose for the sake of argument that human organisms display a considerable degree of indeterminism. Let us suppose in fact that each human organism is such that when the human person associated with that organism (we leave aside the question whether the person and the organism are identical) is trying to decide whether to do A or to do B, there is a physically possible future in which the organism behaves in a way appropriate to a decision to do A and that there is also a physically possible future in which the organism behaves in a way appropriate to a decision to do B. We shall see that this supposition leads to a mystery. We shall see that the indeterminism that seems to be required by free will seems also to destroy free will.

Let us look carefully at the consequences of supposing that human behavior is undetermined. Suppose that Jane is in an agony of indecision; if her deliberations go one way, she will in a moment speak the words, "John, I lied to you about Alice," and if her deliberations go the other way, she will bite her lip and remain silent. We have supposed that there is a physically possible future in which each of these things happens. Given the whole state of the physical world at the present moment, and given the laws of nature, both of these things are possible; either might equally well happen.

Each contemplated action will, of course, have antecedents in Jane's cerebral cortex, for it is in that part of Jane (or of her body) that control over her vocal apparatus resides. Let us make a fanciful assumption about these antecedents, since it will make no real difference to our argument what they are. (It will help us to focus our thoughts if we have some sort of mental picture of what goes on inside Jane at the moment of decision.) Let us suppose that there is a certain current-pulse that is proceeding along one of the neural pathways in Jane's brain and that it is about to come to a fork. And let us suppose that if it goes to the left, she will make her confession, and that if it goes to the right, she will remain silent. And let us suppose that it is undetermined which way the pulse will go when it comes to the fork: even an omniscient being with a complete knowledge of the state of Jane's brain and a complete knowledge of the laws of physics and unlimited powers of calculation could say no more than, "The laws and the present state of her brain would allow the pulse to go either way; consequently, no prediction of what the pulse will do when it comes to the fork is possible; it might go to the left, and it might go to the right, and that's all there is to be said."

Now let us ask: Does Jane have any choice about whether the pulse goes to the left or to the right? If we think about this question for a moment, we shall see that it is very hard to see how she could have any choice about that. Nothing in the way things are at the instant before the pulse makes its "decision" to go one way or the other makes it happen that the pulse goes one way or goes

the other. If it goes to the left, that *just happens*. If it goes to the right, that *just happens*. There is no way for Jane to *influence* the pulse. There is no way for her to *make* it go one way rather than the other. Or, at least, there is no way for her to make it go one way rather than the other and leave the "choice" it makes an undetermined event. If Jane did something to make the pulse go to the left, then, obviously, its going to the left would *not* be an undetermined event. It is a plausible idea that the only way to have a choice about the outcome of a process is to be able to arrange things in ways that will make it inevitable that this or that outcome occur. If this plausible idea is right, then it would seem that there is no way in which anyone could have any choice about the outcome of an indeterministic process. And it seems to follow that if, when one is trying to decide what to do, it is truly undetermined what the outcome of one's deliberations will be, then one could have no choice about that outcome. It is, therefore, far from clear that incompatibilism is a tenable position. The incompatibilist who believes in free will must say this: it is possible, despite the above argument, for one to have a choice about the outcome of an indeterministic process. But how is the argument to be met?

Some incompatibilists attempt to meet this argument by means of an appeal to a special sort of causation. Metaphysicians have disagreed about what kinds of things stand in the cause-and-effect relation. The orthodox, or Humean position, is that—although our idioms may sometimes suggest otherwise—causes and effects are always events. We may say that "Stalin caused" the deaths of millions of people, but when we talk in this way, we are not, in the strictest sense, saying that an *individual* was the cause of certain events. It was, strictly speaking, certain *events* (certain actions of Stalin) that were the cause of certain other events (the millions of deaths). It has been suggested, however, that, although events do indeed cause other events, it is sometimes true that individuals, *persons* or *agents*, cause events. According to this suggestion, it might very well be that an event in Jane's brain—a current-pulse taking the left-hand branch of a neural fork, say—had Jane as its cause. And not some event or change that occurred within Jane, not something Jane *did*, but Jane herself, the person Jane, the agent Jane, the individual thing Jane.

This "type" of causation is usually labeled 'agent-causation', and it is contrasted with 'event-causation', the other "type" of causation, the kind of causation that occurs when one event causes another event. An event is a change in the intrinsic properties of an individual or a change in the way in which certain individuals are related to one another. Event-causation occurs when a change that occurs at a certain time is due to a change that occurred at some earlier time. If there is such a thing as agent-causation, however, some changes are not due to earlier changes but simply to agents.

Let us now return to the question that confronts the incompatibilist who believes in free will: How is it possible for one to have a choice about the outcome of an indeterministic process? Those incompatibilists who appeal to agent-causation answer this question as follows: "A process's having one outcome rather than one of the other outcomes it might have had is an event. For an agent to have a choice about the outcome of a process is for the agent to be able to cause each of the outcomes that process might have. Suppose, for ex-

ample, that Jane's deciding what to do was an indeterministic process and that this process terminated in her deciding to speak, although, since it was indeterministic, the laws of nature and the way things were when the process was initiated were consistent with its terminating in her remaining silent. But suppose that Jane caused the process to terminate in her speaking and that she had been able to cause it to terminate in her being silent. Then she had a choice about the outcome. That is what it is to have had a choice about whether a process terminated in A or B: to have caused it to terminate in one of these two ways, and to have been able to cause it to terminate in the other."

There are two "standard" objections to this sort of answer. They take the form of questions. The first question is, "But what does one add to the assertion that Jane decided to speak when one says that she was the agent-cause of her decision to speak?" The second is, "But what about the event *Jane's becoming the agent-cause of her decision to speak*? According to your position, this event occurred and it was undetermined—for if it were determined by some earlier state of things and the laws of nature, then her decision to speak would have been determined by these same factors. Even if there is such a thing as agent-causation and this event occurred, how could Jane have had any choice about whether it occurred? And if Jane was the agent-cause of her decision to speak and had no choice about whether she was the agent-cause of her decision to speak, then she had no choice about whether to speak or be silent."

These two standard objections have standard replies. The first reply is, "I don't know how to answer that question. But that is because causation is a mystery, and not because there is any *special* mystery about agent-causation. How would you answer the corresponding question about event-causation: What does one add to the assertion that two events occurred in succession when one says that the earlier was the *cause* of the later?" The second reply is, "But Jane did have a choice about which of the two events, *Jane's becoming the agent-cause of her decision to speak* and *Jane's becoming the agent-cause of her decision to remain silent*, would occur. This is because she was the agent-cause of the former and was able to have been the agent-cause of the latter. In any case in which Jane is the agent-cause of an event, she is also the agent-cause of her being the agent-cause of that event, and the agent-cause of her being the agent-cause of her being the agent-cause of that event, and so on 'forever.' Of course, she is not *aware* of being the agent-cause of all these events, but the doctrine of agent-causation does not entail that agents are aware of all of the events of which they are agent-causes."

Perhaps these replies are effective and perhaps not. I reproduce them because they are, as I have said, standard replies to standard objections. I have no clear sense of what is going on in this debate because I do not understand agent-causation. At least I don't think I understand it. To me, the suggestion that an individual thing, as opposed to a *change* in an individual thing, could be the cause of a change is a mystery. I do not intend this as an argument against the *existence* of agent-causation—of some relation between individual things and events that, when it is finally comprehended, will be seen to satisfy the descriptions of "agent-causation" that have been advanced by those who claim to grasp this concept. The world is full of mysteries and of verbal de-

scriptions that seem to some to be nonsense and which later turn out to have been appropriate. ("Curved space! What nonsense! Space is what things that are curved are curved *in*. Space itself can't be curved." And no doubt the phrase 'curved space' *wouldn't* mean anything in particular if it had been made up by, say, a science-fiction writer and had no actual use in science. But the general theory of relativity does imply that it is possible for space to have a feature for which, as it turns out, those who understand the theory all regard 'curved' as an appropriate label.) I am saying only that agent-causation is a mystery and that to explain how it can be that someone can have a choice about the outcome of an indeterministic process by an appeal to agent-causation is to explain a mystery by a mystery.

But now a disquieting possibility suggests itself. Perhaps the explanation of the fact that both compatibilism and incompatibilism seem to lead to mysteries is simply that the concept of free will is self-contradictory. Perhaps free will is, as the incompatibilists say, incompatible with determinism. But perhaps it is also incompatible with indeterminism, owing to the impossibility of anyone's having a choice about the outcome of an indeterministic process. If free will is incompatible with both determinism and indeterminism, then, since either determinism or indeterminism has to be true, free will is impossible. And, of course, what is impossible does not exist. Can we avoid mystery by accepting the non-existence of free will? If we simply say that no one ever has any choice about anything, then we need not reject the No Choice Principle, and we need not suppose that it is possible for a person to have a choice about the outcome of an indeterministic process.

But consider. Suppose that you are trying to decide what to do. And let us suppose that the choice that confronts you is not a trivial one. Let us not suppose that you are trying to decide which of two movies to see or which flavor of ice cream to order. Let us suppose that the matter is one of great importance—great importance to *you*, at any rate. You are, perhaps, trying to decide whether to marry a certain person or whether to risk losing your job by reporting unethical conduct on the part of a superior or whether to sign a "do not resuscitate" order on behalf of a beloved relative who is critically ill. Pick one of these situations and imagine that you are in it. (If you are in fact faced with a non-trivial choice, then you have no need to imagine anything. Think of your own situation.) Consider the two contemplated courses of action. Hold them before your mind's eye, and let your attention pass back and forth between them. Do you really think that you have no choice about which of these courses of action will become actual? Can you really believe that?

Many philosophers have said that although the choice between contemplated future courses of action always seems "open" to them, when they look back on their past decisions, the particular decision that they have made always or almost always seems inevitable once it has been made. I must say that I do not experience this myself, and, even if I did, I should regard it as an open question whether "foresight" or "hindsight" was more to be trusted. (Why should we suppose that hindsight is trustworthy? Maybe there is within us some psychological mechanism that produces the illusion that our past decisions were inevitable in order to enable us more effectively to put these deci-

sions behind us and to spare us endless retrospective agonizing over them. Maybe we have a natural tendency to reinterpret our past decisions in a way that presents them in the best possible light. One can think of lots of not implausible hypotheses that would have the consequence that our present impression that our past decisions were the only possible ones—if we indeed have this impression—is untrustworthy.)

When I myself look at contemplated future courses of action in the way I have described above, I find an irresistible tendency to believe that I have a real choice as to which one will become actual. It may be, of course, that this tendency is the vehicle of illusion. If the concept of free choice were self-contradictory, it might still be the case that a belief in this self-contradictory thing was indispensable to human action. What would it be like to believe, really to *believe*, that in every circumstance only one course of action was open to one?

It can plausibly be argued that it would be impossible under such circumstances ever to try to decide what to do. Suppose, for example, that you are in a certain room that has a single door and that this door is the only possible exit from the room. Suppose that, as you are thinking about whether to leave the room, you hear a click that may or may not have been the sound of the door being locked. You are now in a state of uncertainty about whether the door is locked and are therefore in a state of uncertainty about whether it is possible for you to leave the room. Can you continue to try to decide whether to leave the room? It would seem that you cannot. (Try the experiment of imagining yourself in this situation and seeing whether you can imagine yourself continuing to try to decide whether to leave.) You cannot because you no longer believe that it is possible for you to leave the room. It's not that you believe that it is impossible for you to leave the room. You don't believe that either, for you are in a state of uncertainty about whether it is possible for you to leave. You can, of course, try to decide whether to get up and try the door. But that is something—or at least you probably believe this—that is possible for you. And you can try to decide, conditionally, as it were, whether to leave the room if the door should prove to be unlocked. But that is not the same thing as trying to decide whether to leave the room.

This thought-experiment convinces me that I cannot try to decide whether to do A or B unless I believe that doing A and doing B are both possible for me. And, therefore, I am convinced that I could not try to decide what to do unless I believed that sometimes more than one course of action was open to me. And if I never decided what to do, I should not be a very effective human being. In the state of nature, I should no doubt starve. In a civilized society, I should probably have to be institutionalized. Belief in one's own free will is therefore something that we can hardly do without. It would seem therefore that it would be an evolutionary necessity—at least for rational beings like ourselves—that we believe in free will. And evolutionary necessity has scant respect for such niceties as logical consistency. It is arguable, therefore, that we cannot trust our conviction that we have free will (if, indeed, we do have this conviction). If evolution would force a certain belief on us—by brutally culling out all those of our ancestors who lacked this belief—then the fact that

we hold this belief is no evidence whatever that the belief is true or even logically consistent. (But *aren't* there people who believe that no one has free will, including themselves? Well, there are certainly people who *say* that they believe this, but I suspect that they are not describing their own beliefs correctly. But even if there are people who believe that no one has free will, it does not follow that these people do not believe in free will, for people do have contradictory beliefs. It may be that "on one level"—the abstract and theoretical—certain people believe that no one has free will, although on another level—the concrete and everyday—they believe that people have free will.)

Nevertheless, when all is said and done, I find myself with the belief that sometimes more than one course of action is open to me, and I cannot give it up. (As Dr. Johnson said, "Sir, we know our will is free, and there's an end on't.") And I don't find the least plausibility in the hypothesis that this belief is an illusion. It can sometimes seem attractive to hold the view that free will is an illusion. To think this—or to toy with the idea in a theoretical sort of way—can be attractive to someone who has betrayed a friend or achieved success by spreading vicious rumors. If you had done something of that sort, wouldn't you want to believe that you had no choice, that no other course of action was really open to you? Wouldn't it be an attractive idea that your actions were determined by your genes and your upbringing or even by the way things were thousands or millions of years ago? (Jean-Paul Sartre once remarked that determinism was an endless well of excuses.) And it is immensely attractive to suppose that one is a member of a very small minority that has seen through an illusion that people have been subject to for millennia. The hypothesis has its unattractive aspects too, of course. For one thing, if it rules out blame, it presumably rules out praise on the same grounds. But, however attractive or unattractive it may be, it just seems to be false. If some unimpeachable source—God, say—were to tell me that I didn't have free will, I'd have to regard that piece of information as proof that I didn't understand the World at all. It would be as if an unimpeachable source had told me that consciousness did not exist or that the physical world was an illusion or that self-contradictory statements could be true. I'd have to say, "Well, all right. You *are* an unimpeachable source. But I just don't understand how that could be right."

I conclude that there is no position that one can take on the matter of free will that does not confront its adherents with mystery. I myself prefer the following mystery: I believe that the outcome of our deliberations about what to do is undetermined and that we—in some way that I have no shadow of an understanding of—nevertheless have a choice about the outcome of these deliberations. (And I do not believe that the concept of agent-causation is of the least help in explaining how this could be.)

I believe that if Jane has freely decided to speak then the following must be true: if God were to create a thousand perfect duplicates of Jane as she was an instant before the decision to speak was made and were to place each one in circumstances that perfectly duplicated Jane's circumstances at that instant, some of the duplicates would choose to speak and some of them would choose to remain silent, and there would be no explanation whatever for the

fact that a particular one of the duplicates made whichever of the choices it was that she made. And yet, I believe, Jane had a choice about whether to speak or to remain silent. (It is important not to be misled by words here. From the fact that someone *makes* a choice, it does not follow that that person *has* a choice. If I am locked in a room and do not know that the door is locked, it may be that I *make* a choice *to* stay in the room even though I *have* no choice *about whether* to stay in the room.)

I accept this mystery because it seems to me to be the smallest mystery available. If someone believes that human beings do not have free will, then that person accepts a mystery and in my view it is a greater, deeper mystery than the one I accept. If someone denies the No Choice Principle, then that person accepts a mystery, and in my view it is a greater, deeper mystery than the one I accept. But others may judge the "sizes" of these mysteries differently.

It is important to be aware that we have not said everything there is to be said about the size of the mysteries. The most important topic that we have not discussed in that connection is the relation between free will and morality. In our preliminary discussion of the concept of free will, we said that it was a common opinion that free will was required by morality. If this common opinion is correct, then all moral judgments are false or in some way "out of place" if there is no free will. If that were so, it would greatly aggravate the mystery that confronts those who deny that there is free will. Could it really be true, for example, that those who believe that there is something morally objectionable about racism or child abuse or genocide or serial murder hold a belief that is false or in some other way defective? If an unimpeachable source were to inform me that there was nothing morally objectionable about child abuse, my dominant reaction would be one of horror. But I should also have a negative reaction to this revelation that was more intellectual or theoretical. I should have to say that, if that was so, then I didn't understand the World at all. I should have to say that I just didn't understand how it could *be* that there was nothing morally objectionable about child abuse.

It may not be, however, that those who reject free will must hold that all moral judgments are false or otherwise illusory. The "common opinion" that morality requires free will is not so common as it used to be. When almost all English-speaking philosophers were compatibilists, this opinion was held by almost everyone in the English-speaking philosophical world. It was the common assumption of the compatibilists and the few incompatibilists that there were. Now, however, compatibilism is a less common opinion, owing to the fact that philosophers are coming to realize that compatibilism requires the rejection of the No Choice Principle. Many philosophers are now inclined to reject compatibilism who would previously have accepted it. And because they are also inclined to reject the view that we could have free will in a way that required indeterminism, they are inclined to reject free will altogether. But most of them are not willing to say that morality is an illusion. It has, therefore, become an increasingly widespread view that morality does not after all require free will. It is because of this increasingly popular view that I have not included the thesis that morality is an illusion among the mysteries that must be accepted by those who reject free will. I myself continue to believe that moral-

ity is an illusion if there is no free will, but, since the issues involved in the debate about this question pertain to moral philosophy rather than to metaphysics, I shall not discuss them.

However one may judge the relative "sizes" of the mysteries that confront the adherents of the various positions that one might take on the question of free will, these mysteries exist. The metaphysician's task is to display these mysteries. Each of us must decide, with no further help from the metaphysician, how to respond to the array of mysteries that the metaphysician has placed before us.

Suggestions for Further Reading

Berofsky's *Free Will and Determinism* and Watson's *Free Will* are excellent collections devoted to the problem of free will and determinism. Fischer's more recent *Moral Responsibility* contains much useful material. My own book, *An Essay on Free Will* is a defense of incompatibilism. Large parts of it are accessible to those without formal philosophical training. The central argument of the book is attacked in Lewis's superb article, "Are We Free to Break the Laws?" (rather difficult for those without philosophical training). Dennett's *Elbow Room* is a highly readable (if somewhat idiosyncratic) defense of compatibilism.