

Rational Action, Freedom, and Choice

E. J. Lowe

Is a naturalistic account of rational human action possible? Obviously, we can't answer this question without being told what the questioner means by a 'naturalistic account' of some phenomenon. For most philosophers, however, 'naturalistic' just means 'physicalistic'. For many of these philosophers, a naturalistic account of rational human action would be one which represented human actions as being wholly physical events with wholly physical causes and the rationality of an action could only have something to do with how and by what it was caused. For example, it might be held that every human action is simply a bodily movement of some kind and that such an action qualifies as rational just in case it was caused by the onsets of certain psychological states of the agent whose contents represented such an action as serving the agent's interests in the circumstances in which the agent found himself. For such an account to qualify as wholly naturalistic in the sense now under consideration, it would have to incorporate a physicalistic account of mental representation and of an agent's interests. Perhaps such an account of mental representation could be provided in causal or teleofunctional terms and perhaps such an account of an agent's interests could be provided in terms of evolutionary adaptation. I shall not, however, be pursuing questions about mental representation any further in the present paper, because I think that a physicalist account of rational human action along the lines just proposed inevitably falls at an earlier hurdle, in virtue of endeavouring to explain the occurrence of such actions in wholly causal terms. I shall give my reasons for saying this shortly. Before coming to these reasons, though — which are, of course, reasons for *belief* — I need to say something about the notion of a reason for *action*. Only at the very end of the paper shall I say something about the connection between reasons for belief and reasons for action.

1. Acting for a reason

What is an action? What is a reason to act? And what is it to act for a certain reason? An action is something done by an agent, such as opening a door, raising an arm, uttering a sentence, or imagining a situation. Some actions, such as imagining, are mental actions. Others, such as raising an arm, are physical actions. An ‘agent’, in the sense now intended, is a psychological subject — something that can have thoughts, feelings, desires, intentions, urges and so forth. Persons or selves are psychological subjects with a capacity for reason and self-reflection. And it is with such agents — that is, with beings like ourselves — that I shall be exclusively concerned in what follows.

On the issue of the nature of reasons for action, I espouse what would commonly be called an ‘externalist’ view. A reason for an agent *S* to perform some action *a* consists, typically, in some fact concerning the agent’s circumstances. For example, if *S* is standing near to a dilapidated building, the fact that a loose slate is sliding from the roof and is about to fall on *S*’s head constitutes a reason for *S* to jump out of the way. Of course, *S* will be in no position to act upon that reason and jump out of the way unless *S* is aware of the impending danger. It seems safe to say, then, that only if *S* *believes* that the slate is about to fall on his head will *S* jump out of the way for that reason. But this is very far from saying either that *S*’s *belief* that the slate is falling is his reason, or part of his reason, for jumping or that this belief is a *cause* of his jumping. *That* an agent has a certain belief can, but typically isn’t, a reason for the agent to act in a certain way — but *a belief* as such cannot ever be a reason for action, because it is not even something in the right ontological category to constitute such a reason. And, although a belief can sometimes cause or help to cause some of an agent’s behaviour, a rational action is precisely one that *isn’t* caused by an agent’s beliefs and desires. Rather, it is one that the agent *chooses* to perform ‘in the light of’ his or her beliefs and desires, which is a very different matter.

To expand a little upon these claims, consider first a case in which the fact that an agent has a certain belief constitutes a reason for the agent to act in a certain way. Here I can do no better than adapt an example of Jonathan Dancy’s.¹ The fact that someone has the paranoid belief that he is being pursued by alien beings from Mars is a good reason for that person to visit a psychiatrist — although, of course, it is

unlikely that such a person will be able to act upon this reason, because it is unlikely that he will be aware that his belief is paranoid. By contrast, in our earlier example of the falling slate, it is not the fact that *S* believes that the slate is falling that constitutes a reason for *S* to jump out of the way: rather, it is the fact that the slate is falling — although, once more, *S* must be aware of this fact and so believe, or at least suspect, that the slate is falling if he is to be in a position to act upon that reason. The mistake that is characteristically made by ‘internalists’ is to regard such psychological states of the agent as constituting reasons for the agent to act, rather than the *contents* of such states — that is, those facts or states of affairs whose existence is revealed to the agent in virtue of his being in such psychological states. The mistake is twofold in character. First, it represents a reason for action as being a state of the agent, when in fact it is a state of affairs. Second, it represents a reason for action as always being psychological in nature, when in fact this is only the case in exceptional circumstances, such as that of the paranoid believer. Even in the latter sort of case, the relevant reason for action is a psychological *state of affairs* — the fact that the agent has the paranoid belief — rather than the psychological state of the agent, his paranoid belief.

Next let us consider a case in which an agent’s beliefs do cause, or help to cause, some behaviour of the agent’s. The example of the paranoid believer in persecuting Martians will again serve us well. An agent with such a paranoid belief may well be caused by that belief to act in various seemingly bizarre ways, such as hiding behind bushes when approached by strangers in the park. Now, of course, being pursued by Martians might well constitute a good reason to hide behind bushes and so we are in a better position to understand this agent’s seemingly bizarre behaviour once we know that he has the paranoid belief. But we are not yet in a position to judge that the agent is acting rationally, or for a reason, once we know that he has this belief and that its content provides him with a reason for acting in the way that he does. If his paranoid belief is *causing* him so to act, then he is not *choosing* to act in this way in the light of that belief: he is simply being impelled by the belief and as such is not acting rationally. It is already well known that there is a problem of ‘deviant causal chains’ for those theorists who maintain that to act rationally is to be caused to act in a certain way by beliefs and desires which represent such an action as being in the agent’s interests in the circumstances in which he finds himself. To use

Donald Davidson's own famous example, a climber may be caused to let go of the rope supporting his companion by his strong desire to save himself and his belief that by letting go he would increase his own chances of survival — and yet the climber's behaviour may be unintentional, because it was only that the onset of this belief and desire so unnerved him that his grip slackened.² It has proved notoriously difficult to say, quite generally, *how* beliefs and desires must cause behaviour if the behaviour is to qualify as rational action performed in the light of those beliefs and desires. What I am now contending, however, is that there *can* be no correct account of rational action in such causal terms, because an action's being *caused* by the agent's beliefs and desires is incompatible with its being an action which the agent *chose* to perform in the light of those beliefs and desires — and no action is rational if the agent does not freely choose to perform it. An action can be *in accordance with reason*, or *reasonable*, without being rational. To be rational, it must be done *for* a reason which the agent freely chooses to act upon — or so I claim. For example, jumping out of the way of a falling slate is a reasonable thing to do, but it is not done for a reason and so rational unless the agent chose to jump out of the way in the light of his belief that the slate was falling.

This may seem to turn on its head a famous argument of Davidson's in defence of the doctrine that 'reasons are causes'.³ And so it does. The argument points out that it is not enough, to explain why an agent acted in the way he did, just to point out that the agent was aware of certain reasons for so acting. For the agent might well have been aware of various different reasons for so acting and, indeed, of various reasons for acting in other ways instead. To explain why the agent acted as he did, we need to know which of these reasons was the reason for which he actually did so act. We are supposed to conclude that the reason for which he actually acted was the one that actually *moved* him so to act, that is, the one that *caused* him so to act — which would have to be a psychological state of the agent, such as a certain combination of belief and desire. However, I have just maintained that being *caused* to act by certain of one's beliefs and desires is in fact incompatible with one's acting rationally, even though it may be compatible with one's acting 'reasonably', or 'in accordance with reason'. And the problem of deviant causal chains is grist to my mill here. But this still leaves me needing to answer the question that Davidson's argument poses: which of those reasons for acting of which the agent may have been aware at the time

of acting was the reason for which he actually acted, if there was one? But my answer to this is straightforward: the reason for which the agent acted is the reason which the agent *chose* to act upon. Normally, the agent himself will be able to tell us which reason this was. In situations in which an agent is aware of a variety of reasons for acting in a certain way, and other reasons for acting in other ways, the agent is confronted with a choice as to how to act and on which reason to act. By choosing to act in a certain way in the light of a given reason, the agent *makes* that reason the reason for which he acted on that occasion.⁴ This is something that the agent himself brings about through his mental act of free choice.

2. Objections and replies

I am aware, of course, that this talk of a ‘mental act of free choice’ will ring alarm bells with all right-thinking naturalists. I shall say more about it in due course. Let me say at once, though, that talk of such mental acts raises no spectre of a vicious infinite regress. There is no danger of our being compelled to say that each mental act of choice must itself be the consequence of a prior act of choice. If choices are to play the role in rational action that I have assigned to them, they must be uncaused and so ‘free’ in the libertarian sense of ‘freedom’. *A fortiori*, they cannot be caused by prior choices. Choices can be causes of subsequent events, but they cannot be effects of prior ones. I shall return to this matter shortly. First, there are some other matters to clear up.

A reason for action, I have claimed, is some fact or state of affairs. More particularly, it is a fact or state of affairs which justifies a particular course of action — as, for example, the fact of the slate’s falling justifies the action of jumping out of its way. But two problems are likely to be raised for me here. First, what do I say about an agent who jumps in the *mistaken* belief that a slate is about to fall on his head? Second, how can we say that it is rational, *tout court*, to jump out of the way of a falling slate: mustn’t we say, rather, that it is rational to do this *if* one desires not to be hit by it? In which case, mustn’t we admit, after all, that psychological states of the agent are at least partly constitutive of an agent’s reasons for action?

In answer to the first question, I am inclined to say the following. Whether or not an agent is mistaken in believing that a slate is about to fall on his head, *that* a slate is about to fall on one’s head is, by any standard, a good reason to jump out of

the way. If an agent is to act rationally in a certain way, he must be aware of a reason for so acting. But ‘being aware of a reason’ in this sense doesn’t require that the reason should actually obtain. Here we might usefully deploy Alvin Plantinga’s way of talking about facts and states of affairs.⁵ A state of affairs may or may not actually obtain, but if it does it is a fact and may consequently be represented by a true proposition. Reasons for acting are states of affairs and one must be aware of them if one is to act rationally. If the world is as the agent believes it to be, a reason for acting of which he is aware will be a state of affairs that obtains, and so a fact. If the relevant beliefs of the agent are false, he may still be aware of the same reason for acting and act rationally in the light of that reason, but the reason for which he acts will turn out to be a state of affairs that does not actually obtain.⁶ In that case, the agent might well have done better to have acted differently, but he did not act irrationally. Since reasons are not causes, it does not matter that an agent’s reason for acting in a certain way may on occasion be a state of affairs which does not obtain and hence which is not a fact. Reasons for action, even in the case of agents with mistaken beliefs, remain perfectly objective and, typically, non-psychological in character. (Here I should concede that I am glossing over the further question of what distinguishes ‘good’ from ‘bad’ reasons for action, if indeed such a distinction is a proper one to make. No doubt, for an action to be rational, the reason in the light of which the agent performed that action must at least be one that the agent *judged* to be ‘good’, in the sense that he judged the state of affairs in question to justify acting as he did. But this is a further complication which I do not have time to go into at present.)

Perhaps my answer to the first question will not satisfy everyone. For those who are dissatisfied with it, I have a slightly different answer which they may prefer. (For my own part, I am presently undecided as to which is the better answer.) We could say that if an agent jumps out of the way in the light of a *mistaken* belief that a slate is about to fall on his head, then he does not *really* have a reason to act in that way on that occasion. At the same time, however, we can and should insist that such an agent is not acting *irrationally* merely insofar as his belief is false — for if it had been true, he really would have had a reason to act in precisely the way that he does. On this view of the matter, only *facts* — that is, states of affairs which obtain — can actually *be* reasons for action, but an agent can, nonetheless, act ‘for a reason’ which does not actually obtain and in doing so act rationally in precisely the same sense in which an

unmistaken agent can. Most importantly, we do not have to regard the mistaken agent's *false belief* — a psychological state of the agent — as being his 'reason for action', and thus treat mistaken and unmistaken agents asymmetrically.

So let us turn to the second question raised a moment ago, namely, how a mere state of affairs, such as the falling of a slate, could constitute a reason for acting in a certain way save in conjunction with some appropriate desire on the part of the agent, such as a desire not to be hit by the slate. What I am inclined to say about this is similar to what I said about the agent's beliefs earlier: first, that a desire as such is the wrong sort of thing to constitute a reason for action and, second, that the fact that the agent has a certain desire, while it could constitute a reason, or part of a reason, for the agent to act in a certain way, does not normally, much less necessarily, do so. Here I am again simply following Dancy's lead. However, I am happy — as Dancy may or may not be — to say that some reasons for action manifest themselves to the agent's consciousness as the contents of desires while others do so as the contents of beliefs. In the case of the falling slate, the agent's awareness of this reason for jumping manifests itself as his belief, or at least suspicion, that the slate is falling. But jumping in these circumstances is a reasonable thing to do only for an agent in whose interest it is not to be hit by such a slate. Human beings can be seriously injured or killed by falling slates and it is consequently in their interests not to be hit by them. They will, correspondingly, tend to desire not to be hit by falling slates. The full reason for a human being to jump out of the way of a falling slate is that the slate will very probably injure him and it is in his interest not to be injured. However, although his awareness of this interest will manifest itself, at least in part, as a desire not to be injured, his action of jumping out of the way of the slate will only qualify as a free and so rational action if he *chooses* to jump out of the way in the light of this desire. If the desire merely *causes* him to jump out of the way and the power of choice is not exercised by the agent on this occasion, his behaviour is undoubtedly reasonable, in that it furthers his interests, but it is not an instance of rational action. Or so I want to claim.

Another complication that I should at least touch upon here concerns the distinction between between moral and practical reason. I have been talking so far only about reasons for action which consist in the agent's having certain 'interests' and his being in a certain kind of situation — for instance, in his being in danger of

being hit by a falling slate and its being in his interest not to be injured by it. But sometimes we consider it rational for an agent to act contrary to his own interests, in the furtherance of some moral consideration, such as the requirement to alleviate the suffering of others. I am perfectly happy to endorse this view and to accommodate it by allowing that some reasons for action consist in moral facts or states of affairs. Of course, many naturalists think that the very notion of a ‘moral fact’ is at best obscure and at worst unintelligible — witness John Mackie’s famous argument from ‘queerness’ against such facts.⁷ Not being a naturalist myself, I have no such problem. This is not, of course, to say that nothing more needs to be said about how and why moral facts can exist. However, morality is not my present concern and I shall say no more about it. Many rational actions are purely prudential and have nothing of a moral character about them — and it is with such actions that I am now primarily concerned. What I now want to discuss is the fundamentally metaphysical question of how it is that ‘free choice’, which I take to be an essential ingredient in all rational action, is possible — and how, in particular, we can avoid the conclusion that mental acts of ‘choice’, assuming them to occur at all, would have to be either causally determined by prior events or else merely the random products of chance. For neither of these alternatives is compatible with the view, which I wish to support, that in acting rationally we freely choose between genuine alternatives in the light of the reasons for action of which we are aware.

3. Choice, causation and free agency

Here is the picture of rational action that I wish to endorse. In certain situations, a rational agent is confronted with two or more alternative courses of action, which are genuine alternatives in the sense that, at the time at which the agent must choose between them, no sufficient cause already exists of one of the prospective outcomes. The agent then deliberates, by weighing the reasons for and against the various different alternatives. Finally, the agent makes a decision and chooses to act in one way rather than another. His mental act of choice may then, and normally will, contribute causally to a chain of subsequent events issuing in the performance of the chosen action. So, for example, the agent is attending a lecture and is confronted with a choice between raising his arm to ask the speaker a question or refraining from raising his arm and instead listening to other people’s questions. The agent is aware of

certain considerations favouring the first course of action and other considerations favouring the second. For instance, in favour of the first may be the fact that his question is an important one and no one else is likely to raise it. In favour of the second may be the fact that, being of a shy and nervous disposition, he is not as good as many other people at articulating questions clearly in public. Having deliberated a while, the agent then decides to ask his question, chooses to raise his arm and, by this act of choice initiates a train of events which results in the rising of his arm. We are to suppose that, prior to the agent's act of choice, no causally sufficient condition of the arm's going up when it did already existed. The first moment at which there existed a causally sufficient condition of this event was the moment at which the agent chose to raise his arm. The agent's act of choice was perfectly free, in the sense that it had no antecedent cause whatsoever. Instead of that act of choice occurring, it was equally possible, given the state of the universe prior to that moment, that a different choice should have been made by the agent — a choice to refrain from raising his arm — in which case his arm would not have risen when it later did.

I suspect that many people will agree that this is how they think of themselves as proceeding when confronted with what appears to be a choice between two or more alternative courses of action. But does this picture really make sense? Naturalistic philosophers will almost certainly say no — and even many antinaturalistic philosophers may be inclined to agree with them. Some will argue as follows. Suppose the agent chooses to raise his arm and his arm duly goes up. We need to ask *why* the agent chose to raise his arm. There must surely be some explanation as to why the agent chose as he did and the most plausible explanation is that the agent's mental state immediately prior to the act of choice caused it to occur. So, it may be said, at some stage in the agent's deliberations he formed the belief that, all things considered, asking a question was the sensible thing to do and this belief, together with his desire to do what he considered sensible, caused him to choose to act in the way he did, by raising his arm. However, now we are back to the picture of rational action according to which a rational action is one that is not freely chosen in the light of the agent's beliefs and desires but *caused* by the agent's beliefs and desires. And this is precisely the picture that I was at pains to reject earlier on. I think it is a picture that is radically false to the phenomenology of rational action and, worse still, impossible to square with any conception of rationality that we can really make sense

of. That is to say, when we take ourselves to be acting rationally it never *seems* to us that we are being caused to act in the ways we do by our beliefs and desires, and once we suppose that, on a given occasion, we were in fact caused to act in a certain way by our beliefs and desires, we find ourselves obliged to withdraw any claim to have acted rationally on that occasion.

Let us go back to the question that precipitated these difficulties. The question was *why* the agent chose to raise his arm. I suggested that what prompts this question is the thought that there must surely be some *explanation* as to why the agent chose as he did — and I then suggested that, at least for many philosophers, the most plausible explanation is that the agent's mental state immediately prior to the act of choice caused it to occur. But suppose the question is addressed to the agent and he is asked 'Why did you chose to raise your arm?'. The agent will no doubt answer that he chose to raise his arm in order to ask a question. If pressed a little further, he will say, perhaps, that he chose to raise his arm because he had an important question to ask which no one else was likely to raise. In other words, he will cite certain reasons that he was aware of for raising his arm. Pressed still further, he may admit that he was also aware of certain reasons for refraining from raising his arm. 'So', we may ask him, 'why did you chose to act on the first set of reasons and not on the second?'. He may answer, 'Well, after some deliberation I came to the conclusion that the first set of reasons was better than the second, so I chose to act on the first.' For most ordinary folk, this is where the questioning would come to an end. But philosophers, like little children, sometimes don't know when to stop asking 'why?'. They may now ask, 'But why did you choose to act on the reasons that you thought were the better ones?'. Here our agent, if he is not a philosopher, is likely to return a blank stare of bewilderment. If he says anything, it is likely to be sarcastic or rhetorical, such as, 'Well, don't *you* usually choose to act on what you consider to be the best reasons?'. Our philosophical interrogator may now pounce and say, 'So, you couldn't help choosing as you did, once you came to the conclusion that one set of reasons was better than the other — in other words, your coming to that conclusion *caused* you to choose as you did'. But our agent is likely to resist this suggestion very forcefully and reply, 'No, of course I *could* have chosen to act otherwise and nothing *made* me choose to act as I did — though if I had chosen to act otherwise, it would have been

against my better judgement, and sometimes I have done precisely that and later regretted it’.

4. Choice and chance

At this point, our philosophical interrogator, seeing that he is not going to get the agent to admit that anything *caused* him to choose as he did, may take another tack and try to persuade the agent that if his choice was genuinely uncaused, then it must have been somehow arbitrary — a mere chance event — and as such the very antithesis anything rational. Moreover, if acts of choice are mere chance events, then how can any agent be said to have an *ability* to choose to act in one way rather than another? For to possess such an ability the agent must surely have some sort of *control* over how he chooses and no one can, almost by definition, have any sort of control over what happens merely by chance.

One way of pursuing this line of argument is through a thought experiment suggested by Peter van Inwagen.⁸ We are to suppose that Alice is faced with a choice between lying and telling the truth on a given occasion and, after much deliberation, chooses to tell the truth. The libertarian will want to represent this choice as being genuinely free, in the sense of being uncaused. This means that, according to the libertarian, the state of the universe prior to Alice’s act of choice contained no causally sufficient condition of that act and, up until the moment of choice, there were possible futures in which Alice tells the truth and also possible futures in which she lies. Suppose, then, that all of the circumstances leading up to the moment of Alice’s choice were to be replicated or ‘replayed’ a number of times — say, one thousand times. Since we are supposing that Alice’s choice in the original ‘play’ of these circumstances was causally undetermined by preceding states of affairs, we have to suppose that Alice’s choice is likewise causally undetermined by preceding states of affairs in each of the replays. So, it seems, we must suppose that in some of the replays she chooses as she did originally and tells the truth and in the others she chooses differently and tells a lie. It seems that there will be some specific number of replays in which she chooses to tell the truth — say, 513 — and in the remaining 487 replays she chooses to tell a lie. But it was arbitrary to pick one thousand as the number of replays. We can generalize and say that there will be some particular proportion of all the possible replays of Alice’s situation in which she chooses to tell

the truth. This proportion would have to be less than unity, for to suggest that Alice would have chosen to tell the truth in *every* possible replay of the situation is surely to imply that her choice was, after all, determined by preceding states of affairs. But if the proportion is less than unity — say, it is one half — then what this appears to signify is that Alice's choice in the original situation was simply a chance event whose objective probability of occurrence was 0.5, or 50%.

The first thing I want to say about this argument is that, even if we can make complete sense of the notion of a 'replay' of the circumstances preceding Alice's act of choice and even if we can agree that, given that her choice was causally undetermined, it follows that in some replays she would choose to tell the truth while in others she would choose to tell a lie, it doesn't follow that there must be some determinate proportion of all possible replays in which she chooses to tell the truth. In other words — and avoiding the perhaps questionable notion of a 'replay' of Alice's situation — the fact that Alice might have chosen otherwise than she actually did does not imply that there was a certain numerically precise objective probability of her choosing as she did in the one 'play' that actually did happen. Consequently, it is far from clear that we are entitled to regard her choice as a 'chance event'. Compare Alice's situation with that which occurs in a genuine game of chance, such as a dice game. Let us suppose that on a certain occasion a perfectly fair die is thrown and it lands with the six uppermost. And let us suppose that its so landing is a genuinely indeterministic event, like the spontaneous decay of a radium atom — even though there is reason to doubt that this is really so. If we now contemplate all possible 'replays' of this situation, we do indeed have some reason to suppose that a definite proportion of them — one sixth — will result in the die landing with the six uppermost, and hence that there was an objective probability or chance of one in six of this happening when the die was actually thrown. Our reason to think so is obviously based on symmetry considerations concerning the structure of the die. However, nothing remotely like this reason is able to support a similar conclusion in Alice's case. It may be objected that we likewise lack any such reason to assign a numerical degree of chance to an indeterministic atomic event, such as the spontaneous decay of a radium atom, and yet numerical probabilities are associated with such events. However, this is because large amounts of statistical data enable us to assign 'half-lives' to radioactive isotopes. Nothing similar even to this is actually

available in the case of Alice. By being invited to contemplate a large number of ‘replays’ of Alice’s situation, we are being invited to suppose that statistical data of this sort *is* in principle available concerning Alice’s situation, albeit data that is not all drawn from the *actual* world, as in the case of the data which enables us to assign half-lives to radioactive isotopes: rather, the ‘data’ in Alice’s case is, as it were, distributed across all the possible worlds compatible with the history of actual events up until the moment of Alice’s choice. I suppose that there are some conceptions of possible worlds that would support the idea that such ‘data’ really exists, but for my own part I see no reason to think that it does. I see no reason to think that even God could assign an objective numerical probability or degree of chance to Alice’s choosing to tell the truth.

The next thing I want to say about the Alice thought experiment is this. Alice, we have said, deliberated at some length before deciding how to act and choosing accordingly. The way in which the thought experiment is set up requires us to suppose that in all ‘replays’ of her situation her deliberations proceed in exactly the same way that they actually did, but that at the end of some of these replays she still chooses to tell the truth while at the end of others she chooses instead to tell a lie. However, if Alice deliberates and chooses rationally, at the end of her deliberations she will have formed a judgement as to whether the reasons in favour of telling the truth are better than those in favour of telling a lie. Now, certainly, if Alice’s choice was genuinely undetermined, it was possible for her to have chosen to tell a lie in spite of having formed such a judgement. On the other hand, there is a clear sense in which, in the light of that judgement, Alice would have been acting irrationally in telling a lie. A libertarian should happily accept that our freedom to choose is a freedom even to choose irrationally — and we all recognize, I think, that sometimes we do make irrational choices. So-called ‘akratic’ action involves precisely this, it would seem, for such an action is one that the agent chooses to perform ‘against his better judgement’, as we say. At the same time, little sense can be made of the notion of a rational agent who frequently, or as often as not, chooses to act irrationally as to act rationally.

We have to distinguish clearly between two different conceptions of how Alice might have chosen differently. She might have chosen differently even if the result of her deliberations had been the same, that is, even if she had formed the same

judgement about which action had better reasons in favour of it. A libertarian must certainly say this. But another thing that can and should be said by the libertarian is that Alice might have deliberated differently and formed a different judgement as to which action had better reasons in favour of it, in which case, too, she might have chosen to act differently — only this time she would not have chosen ‘akratically’. The reason why the libertarian should say this is that deliberation itself should be seen as a process which involves choice. Deliberation isn’t simply a matter of the agent’s being confronted with ready-made reasons in favour of doing this or that action and then having to weigh those reasons against one another. Deliberation involves the active seeking of reasons for or against carrying out one of a range of alternative possible actions. Consequently, it involves the agent in making choices about what sources of evidence to consult or think about with a view to revealing reasons for or against acting in a certain way. We can choose not only how to move our bodies but also how to direct our thoughts, and the latter kind of choice is essentially involved in all processes of deliberation. But if this is what deliberation really involves and if choice really is causally undetermined in the way the libertarian maintains, then the libertarian should not meekly accept the terms in which the thought experiment involving Alice is described. For the libertarian should reject the idea that one could, even in principle, set up circumstances at any time prior to Alice’s act of choice in such a fashion that she was bound to go through the same process of deliberation as she actually did. This is not to say that the libertarian should redescribe Alice’s episode of rational deliberation and action simply as a punctuated sequence of discrete indeterministic choices separated by deterministically evolving mental processes — first, choices about how to direct her thoughts in gathering and evaluating reasons for and against truth-telling and finally a choice, in the light of those reasons, whether to tell the truth or a lie. Rather, the whole train of deliberation which culminates in the final choice should be seen as a continuous indeterministic process, capable of evolving differently at every moment at which it is going on.⁹ Seen in this light there is, I think, no intelligible way of representing what is going on with Alice as having any numerical degree of objective probability or chance associated with it, even setting aside my earlier objections to this notion.

The final thing I want to say about the Alice thought experiment is that the choice that Alice actually made was made in the light of reasons of which she was aware favouring the course of action that she chose — and that if she had chosen differently, whether or not after a different process of deliberation, she would still have chosen as she did in the light of reasons of which she was aware. Choice, by its very nature, is never exercised ‘blindly’, but is always informed by or responsive to reasons for action. In this respect, it is utterly unlike any mere chance event, such as the fall of a die or the decay of a radium atom. It doesn’t, I believe, make sense to think of an agent as *choosing* between two or more possible courses of action without being aware of any consideration whatever in favour of carrying out any of them. This is not to deny that an agent may find himself faced with a choice between two courses of action for which equally good reasons seem to present themselves — a ‘Buridan’s ass’ situation — so that he can do nothing but make an arbitrary choice in favour of one of them. Nor is it to deny that we sometimes choose to do things ‘just because we felt like doing them’ — as, for example, when we choose to kick a small stone lying in our path as we stroll along. Even in this sort of case, our choice is made in the light of the consideration that no harm is likely to come from kicking the stone and that we shall get some small degree of pleasure or satisfaction from kicking it.

5. Choice, agency and control

But what about the notion that rational action requires that the agent be able to exercise *control* over what he does? Is the libertarian notion of choice at odds with this idea? That it is at odds with the idea of control is, of course, one thing that was meant to be shown by the thought experiment involving Alice. Alice’s case was supposed to show this by showing that a free choice, on the libertarian conception of what such a choice involves, can only be a mere chance event — and, almost by definition, one cannot have control over a mere chance event. But even if we have succeeded in defending libertarianism against the charge that free choices as it conceives of them must be mere chance events, this doesn’t of itself serve to explain how, or in what sense, according to libertarianism, an agent has control over what he chooses to do. However, the idea that what is needed is an account of how the agent has control over how he exercises his power of free choice is perhaps a confused one. According to libertarianism, it is precisely because we have a power of choice which

we can exercise freely — that is, one whose exercises are not determined by prior events — that we have control over our actions. If our choices were causally determined for us by prior events, then indeed we would *not* have genuine control over our actions. In fact, our ‘choices’ would not really be *choices* at all, in any serious sense. Our choices, I have said, are necessarily informed by or responsive to reasons for action, but are none the less necessarily free, in the sense of being undetermined. The notion of free choice exercised in the light of reason provides us with the very paradigm of what it is to be ‘in control’ of our actions. To demand that we somehow be ‘in control’ of *our choices* seems either superfluous or incoherent. On the one hand, freely exercising our power of choice precisely *is* being ‘in control’ and as such needs no *further* exercise of ‘control’. On the other hand, what is *in* control cannot, for that very reason, *be* controlled in its exercise of control, if this is understood as involving some further agency. If I say to somebody to whom I assign some special responsibility, ‘OK, you are in control now’, I cannot consistently then go on to say, ‘But remember that I shall be controlling your every move’.

To this it may be objected that so-called ‘Frankfurt-style cases’ suggest otherwise.¹⁰ Couldn’t it, in principle, be the case that some mad scientist is monitoring my brain to see what choices I make, allowing only those choices to take effect which are compatible with what he decides should happen — but that it just so happens that everything I choose to do coincides with his plans? In that case, isn’t there a sense in which I am controlling what I do, by exercising my power of choice, even though he has control over me because he would prevent me from doing anything incompatible with his plans? Notice, however, that even if this thought experiment makes sense — which may be questionable — there is no suggestion that the mad scientist has control over my power of choice, only control over its effects. He can make sure that if I choose to do something incompatible with his plans, my choice will be ineffective. But nothing has been said to suggest that he can make sure how I exercise my power of choice — what choices I make. Nor does it make sense that he could do this. For him to have, *per impossibile*, control over how I exercise my power of choice would be for him to *deprive* me of that power, so that this is, logically, not a sort of control that he can have over me. If he had such control over me that I no longer had any power of choice, *he* would now be the only one of us capable of possessing such a power: there would not be two distinct powers of choice,

mine and his, with his the dominant power. He would be in control of me, but not in control of my power of choice, for I would have none. By the same token then, it would make no sense for *me* to be in control of my own power of choice, for then I would have to have two such powers, one dominating the other — and yet the supposedly dominated power of choice would simply be *extinguished*, not controlled, by the supposedly dominating power. This, I think, demonstrates that it is incoherent to demand of the libertarian that he provide an account of how we have ‘control’ over our power of choice. Having a power of choice gives us all the control we could ever have or need.

6. The pragmatic inconsistency of determinism

The final issue I want to address is this. Having argued that our intuitive self-conception is perfectly coherent in representing ourselves as beings whose rationality in action consists in a power of free choice exercised in the light of reason, the question may still seem to remain as to whether this picture of ourselves is true or false. We may have established that it *could* be true, but could it none the less be the case that it is in fact false and that in painting this picture we are deluding ourselves, perhaps in a way which we find psychologically unavoidable? To those who pose such a question I wish to pose the following question in return. Can we make sense of the thought that we might be confronted either with evidence of an empirical nature, or arguments of a philosophical or logical character, which would be rationally compelling and speak in favour of the falsehood of the picture that has been painted? I do not believe that we can, because I do not believe that we can separate rationality of belief from rationality in action. To form a rational belief that we are not free to act in the light of reason, we should have to exercise rationality in action at the very least in directing our thoughts towards putative sources of evidence for that belief and in evaluating the deliverances of those sources of evidence, or in considering and evaluating arguments in favour of that belief. If we are not free to act in the light of reason, then we are not free to deploy our judgement in the light of reason in seeking out and assessing evidence and arguments for or against this or that belief. If, lacking freedom of rational action, we were to acquire the belief that we lack that freedom *not* through the free direction of our thoughts and the free use of our power of judgement, but rather as a consequence of prior causes determining the contents of our beliefs,

then we would not have acquired that belief rationally and would not be rationally justified in holding it. There is a perfectly good sense, then, in which we simply *cannot* rationally believe that we lack freedom of rational action. If we came to believe that we lack that freedom, it would either be a false belief that we had acquired through a faulty exercise of our rationality, or else it would be a true belief which we did not hold rationally because, in virtue of its truth, we would not be rational beings. What cannot be the case is that we should hold the belief rationally and the belief be true. Nothing, therefore, can rationally commend the belief and reason demands that we dismiss it as false.

Notes

¹ See Jonathan Dancy, *Practical Reality* (Oxford: Oxford University Press, 2000), p. 125.

² See Donald Davidson, 'Freedom to Act', in his *Essays on Actions and Events* (Oxford: Clarendon Press, 1980). The example appears on p. 79.

³ See Donald Davidson, 'Actions, Reasons, and Causes', in his *Essays on Actions and Events*.

⁴ Compare John R. Searle, *Rationality in Action* (Cambridge, MA: MIT Press, 2001), p. 16.

⁵ See Alvin Plantinga, *The Nature of Necessity* (Oxford: Clarendon Press, 1974), pp. 44ff.

⁶ Compare Jonathan Dancy, *Practical Reality*, pp. 131ff.

⁷ See J. L. Mackie, *Ethics: Inventing Right and Wrong* (Harmondsworth: Penguin Books, 1977), pp. 38ff.

⁸ See Peter van Inwagen, 'Free Will Remains a Mystery', in Robert Kane (ed.), *The Oxford Handbook of Free Will* (Oxford: Oxford University Press, 2002).

⁹ I am grateful to Storrs McCall for insight on this point. See further Storrs McCall and E. J. Lowe, 'Indeterminist Free Will', *Philosophy and Phenomenological Research*, forthcoming.

¹⁰ See Harry G. Frankfurt, 'Alternate Possibilities and Moral Responsibility', in his *The Importance of What We Care About: Philosophical Essays* (Cambridge: Cambridge University Press, 1988).