

How Can I Possibly Be Free? A Critique of Neurodeterminism.

The argument that free will is incompatible with the fact that we are supposed to operate in a causally closed material world long antedates the specific claim that neuroscience has recently revealed our freedom to be an illusion. However, I am going to spend quite a bit of my talk on neurodeterminism. This is in part because it is so much in fashion that it has actually reached the popular press.

Slide Time and again, we are told that neuroscientists have demonstrated that we do not consciously will our seemingly voluntary actions; that they have a cerebral or evolutionary origin; that it is our biology that is calling the shots. In short, that the discoveries of neuroscience have added weight to the arguments of traditional determinism; that they have demonstrated that we are either not as free as we thought or that we are not free at all; and that, thanks to brain scientists, we now know to be true, on the basis of empirical evidence, what hitherto philosophers and others only feared might be true.

There is something dodgy, of course, about the claim that empirical science can address essentially metaphysical questions such as whether or not human freedom is real. At any rate, if the arguments for determinism were sound, then we would require no data to support them. Be that as it may, some biologists think that what they have discovered about brain activity further supports the case for denying that we have free will. But my focus on *neurodeterminism* is also motivated by this: that a critical examination of neurodeterminism, highlights some of the erroneous assumptions that lie inside determinism *tout court* and reveal what is often allowed to through on the nod, as if it could not be challenged.

Slide Let me begin with a brief statement of the territory I am going to cover. I shall first outline the general case for deterministic universe and the assumption that this is incompatible with free will. This, I suspect, is the most superfluous part of my talk. I will then summarise the arguments behind neurodeterminism. Next I will examine some of the experiments that have caused so much excitement. I will then challenge how those experiments have been interpreted. This will lead to a critique of the conception of action implicit in neurodeterminism and determinism more broadly construed. **Slide** This will open the way to a wider critique of the determinist notion that our actions are determined by material causes that ultimately lie outside of our control. Against this I will argue that our actions are not only expressive of what we are but that they truly originate with us – that, in the case of my actions, I am where the buck, or sufficient of the buck, starts. The key to this argument lies in acknowledging the reality and nature of first-person being. **Slide** Next I will defend the claim, necessary to uphold a belief in free will, that we are able to operate on the law-bound material world as if from the outside. I will define the nature of that outside, arguing that it is made of ‘Thatter’ rather than matter’. I shall remind you, probably unnecessarily that, collectively, humans have deflected the course of nature. I shall end with a reassurance that my emphasis on our collective power and the extent to which we support each other’s ability to exercise free will does not undermine our individual freedom.

Slide Let me briefly waste your time by reminding you of the central case for determinism. **Slide** The most obvious is that every one of our actions is a physical event. Every physical event has a cause and that cause will in turn have causes. Eventually we shall arrive at causes that lie outside of our control and ultimately have nothing to do with us – for example events that happened before we were born. So the ultimate basis for our actions lies outside of us. **Slide** A complementary argument is that physical events are subject to the laws of nature which are, as David Lewis, by definition unbreakable. We cannot therefore deflect the course of events: whatever we do, whatever we think we have brought about, was going to happen anyway.

Slide At least three assumptions are central to the determinist case: a) actions (including their motivations) are naturally delimited events in the causal nexus; b) the causal nexus unfolds in accordance with the laws of nature; c) that nothing can happen, come to pass or brought about outside of the causal nexus. The examination of the neuroscientific case for determinism will enable us to subject these assumptions to critical examination. For what neuroscience brings to the party is not so much data – though, of course people believe it does, hence the excitement – but a set of inherited assumptions derived from the biological framework, and the materialist ontology embedded within it, from which it approaches the person, the mind, consciousness.

Slide Neurodeterminism is of course wedded to the idea that our minds are our brains. Our brains are evolved organs designed, as are all organs, by natural selection to maximize the replicative ability of the genes whose tool the

brain is. We are largely unaware of this. For many this means that we are acting out a biological script at odds with the humanist story we tell ourselves about ourselves as conscious agents. **Slide** More important is the metaphysical or ontological assumption captured by Daniel Dennett:

There is only one sort of stuff, namely *matter* – the physical stuff of physics, chemistry and physiology – and the mind is somehow nothing but a physical phenomenon. In short, the mind is the brain... We can (in principle!) account for every mental phenomenon using the same physical principles, laws and raw materials that suffice to explain radioactivity, continental drift, photosynthesis, reproduction, nutrition and growth.³³

If people are their brains (as we are repeatedly told), they are identified with a piece of matter and this, like all other pieces of matter, is subject to, and cannot escape from, the laws of material nature. Everything that happens in our brains is the product of material events that impinge on them and the events that result from brain activity – notably our actions – are wired into the endless causal net, extending from the Big Bang to the Big Crunch, that is the history of the material universe. Minds and persons are embedded in the physical world. Our destiny, like that of pebbles and waterfalls, is to be pre-destined.

The eminent neurophysiologist Professor Colin Blakemore expressed this view with admirable lucidity in his Reith Lectures *The Mechanics of the Mind*:

The human brain is a machine which alone accounts for all our actions, our most private thoughts, our beliefs...All our actions are products of the activity of our brains. It makes no sense (in scientific terms) to try to distinguish sharply between acts that result from conscious attention and those that result from our reflexes or are caused by disease or damage to the brain.

If we are identical with our brains, or certain neural discharges in them, we must be just as unfree when we are writing a textbook about the management of seizures, or giving Reith lectures on the brain as when we ourselves are in the grip of a seizure: it 'makes no sense in *neuroscientific terms*' to distinguish between these things.

Other writers are not so radical, or so consistent, as Colin Blakemore. They want only to tone down what they feel to be our exaggerated sense of our own autonomy by picking away empirically at what we traditionally regard as freely chosen behaviour. Many psychologists have taken especial pleasure in demonstrating how our decisions are often influenced by stimuli of which we are unaware and that we act for reasons other than those which we believe drive our actions. In general, it seems as if our conscious feelings are less important than we thought they were. **Slide** What is more, neuroscientists have made systematic observations that appear to them to demonstrate that we do not fully *will* our actions; indeed, as Daniel Wegner argues, the only connexion between willing and acting is that both come from the same unconscious source.

These observations have been given further apparent authority by experiments using fMRI scanning and other methods of directly observing brain activity when people are carrying out supposedly voluntary actions. The pattern of activity seen on the scan often, so it is argued, reveals that more is going on than the actor realizes and that the (unconsciously) 'emotional' brain is frequently more engaged than the 'rational' brain.

Slide There isn't time to address the vast literature and I shall not deal properly with four problems with this view: the methodological problems with brain scanning technology upon which much store has been placed; the fundamental confusions of the claim that neuro-science support the mind-brain identity theory; the assumption implicit in the quote from Colin Blakemore that, if neuroscience can't see it, then it doesn't exist; the confusion between biological origins and cultural consequences. **Slide** All of these underlying the problems are dealt with in pitiless detail in my *Aping Mankind* which is out next month.

Let me, however, illustrate the problems with the apparent empirical support for neurodeterminism with two connected studies, carried out a quarter of a century apart, which have been very widely cited and discussed within neuroscience and philosophy and have attracted the attention of the popular press and the wider public, seeming to challenge the notion that we are in charge of our actions. They will illustrate most of the points of philosophical interest that I wish to make.

Slide The first is a famous set of experiments, carried out by the neurophysiologist Benjamin Libet in the 1980s and repeated and refined many times since then, which seem to show that our brain makes decisions to act before our conscious mind is aware of them, so they are not really *our* decisions at all. The neuroscientist Patrick Haggard

described the paper in which they were first described as 'one of the philosophically most challenging.. in modern scientific psychology' . So what did Libet do and what did he find?

In a typical experiment, Libet's subjects are instructed to make a simple movement – to bend their right wrist or the fingers of their right hand - in their own time. Using an EEG, the experimenter records a particular activity in the brain that indicates a readiness to move. This so-called 'readiness potential' (RP) is seen in the part of the cerebral cortex most closely associated with voluntary movement. The RP occurs about half a second before activity in the relevant muscles of the arm or hand, as recorded by an electromyogram, because it takes time for the neural activity in the cortex to translate into events in the relevant muscles. Nothing worrying there. But Libet made another observation that seemed to raise serious questions. He asked his subjects to recall the position a spot revolving round a clock face in order to determine the time when they were first aware of their urge or intention to make a movement. To his surprise, he found that the RP occurred a consistent third of a second *before* the time at which the subjects reported being aware of a decision to move. Libet concluded from this that the *brain* (not the subject or the person) 'decided' to initiate or at least to prepare to initiate the act before there was any reportable subjective awareness of a decision having been made. Put more simply, the cerebral causes of our actions seem to occur *before* our conscious awareness of deciding to perform them.

These findings are open to a range of interpretations, as we shall see, but they cannot be dismissed as mere artefacts of the method of recording, **Slide** though as has recently been shown we infer rather than perceive the moment we decided to act. Nor can the gap between the electrical signal of the initiation of action, the RP, and the awareness of the intention to perform the action be explained away as the interval between forming an intention and being sufficiently reflectively aware of the intention to allocate it to a particular time. This has been demonstrated rather dramatically by more recent work, this time using fMRI.

Slide Chung Siong Soon and colleagues carried out studies in which a succession of letters were displayed on a screen. Subjects were asked to press a left or a right button at a moment of their own choosing and to note the letter which was being displayed at the time they felt that they were making a decision to press the button. The letter was a time marker. Two regions that lit up in the brain predicted the subject's choice of left or right button. Remarkably, the regions in question (in the part of the cerebral cortex associated with voluntary movement) lit up *a full 5 seconds* before the individual was aware of having made a choice. Moreover, there were other areas in the frontal cortex, traditionally ascribed executive powers, that were active no less than *seven* seconds before awareness of the decision. If the delay in the response of the scanner detecting the activity was accounted for, the interval increased to *ten* seconds. Such a delay could not be due to the subject mistiming the intention to move - a possible explanation for Libet's original findings, as it is somewhat tricky to time one's own decisions. The authors concluded that there is a network of high level control areas 'that begins to prepare an upcoming decision long before it enters awareness'. It looks like we don't know what we are doing until we have found that we have done it.

Slide Libet's original interpretation of his own experiments was that they demonstrated that we do not have free will: the brain 'decides' to move, the brain 'initiates' movement. As Libet put it in a more recent paper: 'If the "act now" process is initiated unconsciously, then the conscious free will is not doing it' ¹² . We do, however, have 'free *won't*': we can inhibit movements that are initiated by the brain. We don't quite initiate voluntary processes; rather we 'select and control them', either by permitting the movement that arises out of an unconsciously initiated process or 'by vetoing progress to actual motor activation'. This has been expressed as our ability to 'rubber stamp' decisions that have already been made by neural networks. It is, however, not very clear why the decisions should require rubber stamping. In the personless world of neuroscience, it makes no more sense for us to 'rubber stamp' the decisions of our brain than for a falling pebble to endorse the gravitational field.

Slide To see just how shaky this conclusion is, we need not only to look at the action Libet's subjects were asked to perform but also to fill in some of the context in which they performed it. The action was the simplest imaginable: a flexing of the wrist; a mere movement. That movement was itself only a minute part of a long sequence of movements amounting to a large-scale action which could be described as 'taking part in Dr. Libet's experiment'. This large-scale action began at least as far back as getting up in the morning to visit Dr. Libet's laboratory (after, perhaps, setting the alarm to make sure one was not late); involved consenting to take part in an experiment whose nature and purpose and safety was fully understood; and required (among many other things) listening to and understanding and agreeing to the instructions that were received - and *then* deciding to flex the wrist. In other words, the immediate prior intention, the psychological event timed by Libet, was not the whole story of the action, only a tiny part of it. It was preceded by many others that were minutes, hours, perhaps days, before the action. The real story is not just the flexing of the wrist but

one of a sustained and complex resolve being maintained over a very long time. This includes many large items of behaviour - getting on and off buses, looking for the laboratory, cancelling or declining other commitments so as to be free to attend the lab, and so on – that have many thousands of motor components.

Once this is appreciated, then the temporal relation between the last step, the wrist flexing, and the Readiness Potential seen in the lab becomes unimportant. The decision to participate in the experiment, which alone gave the wrist flexion its meaning, began not milliseconds, seconds, or minutes, but hours before the wrist was flexed. Perhaps weeks, when the person decided to become a subject in the experiment. The flexing of the wrist is just the last component of this action called 'taking part in Dr. Libet's experiment' which would itself be part of a greater intentional whole, such as 'wanting to please Dr. Libet' or 'wanting to help those clever scientists understand the brain as it might one day help doctors to treat my child's brain injury more effectively'.

It now seems less disturbing (or less exciting, according to your taste) that the Readiness Potential preceded the intention to make a movement by a mere 300 to 450 milliseconds, or the brain activity seen on the scan in Soon's experiment was up to 10 seconds in advance of the intention. The specific intention to flex the wrist belongs to a much wider *field* of intention which has temporal depth and 'existential' extensity and is connected with great swathes of the acting individual's self-world (including her know-how, know-that, motives, principles etc). As Tim Crane expressed it, our actions are interconnected, as are intentions, decisions and plans. The fact that the decisions in the Libet experiment seem to follow the actions is also irrelevant, Crane argues, because our actions unfold without there being explicit decisions – except broad brush ones – at every node. When I am walking to the pub to meet you, there isn't a separate decision corresponding to every one of the hundreds of steps I take to get there.

Slide Libet's experiment illustrates how the (neuro)determinist case against freedom is rooted in a very distorted conception of what constitutes an action in everyday life and it shows us what is suspect about determinism more broadly construed. If you want to make voluntary actions seem involuntary the first thing to do is to strip away their context – the self from which they originate, the nexus of meanings that is the world to which they are addressed - and then effectively break them down into their physical elements. This gets you well on the way to eliminating the difference between a twitch and a deliberate action; or between, say, my involuntarily taking part in the experiment (having been carried to the lab in a coma and woken up simply to move my wrist) and my participating in it because I want to help those clever scientists. It is possible to take this denaturing of actions even further. I can, for example, break up the process of giving this lecture into physiological events such as the formation and rupture of cross-bridges in the fibres of my hand muscles. Now it is perfectly obvious that 'I' cannot *do* this. I would not know how to make a muscle cross-bridge if I tried. But it does not follow that I am not giving this lecture freely or that I am not really intending to give it or that it has no relation to my intentions during the days when I wrote it. All that follows is that frameless atoms of actions cannot be specifically intended. But actions are not made of contextless atomic movements. They are indissoluble synthetic wholes, regulated by higher-order intentions that issue not from bodies but from persons. It is this that embeds actions in a first-person reality that is not reducible to third-person, or no person, material events. Let us develop this a bit further.

The intention of the person who is asked to flex her wrist – in life, as in Libet's lab – is aimed not at the movement itself but at the *goal* – or the nested goals - of the movement: to do as Dr Libet requested, to cooperate with the experiment, to help to advance science. Her participation in the experiment originates in a huge space of possibility, a human world, to which a self is addressed, the theatre of its activity, which has been fashioned out of the pooling of our consciousness through shared attention. In the case of Libet's experiments, that world – upon which I have an individual 'take' and through which I make individual tracks – includes the institution of science, my understanding of it, and my attitude towards scientific research, which is interested and sympathetic enough for me to be willing to give up my time to participate in experiments that make sense to me and look as if they may advance our knowledge. It is my self-world which provides the framework and theatre and rationale of the action of flexing my wrist; a boundless hinter-land of meanings that has many layers before it reaches something as simple as a biological or material cause. This is not a delimited portion of the causal nexus, plugged into the remainder of the causal nexus. It is scarcely surprising therefore that we cannot find free will in this isolated movement in a laboratory, if we treat it as an isolated movement, and which treats actions as if they were material events with simple proximate causes and simple proximate effects. The locus of free will is a field of intention, rooted in the self and its world that extends beyond the laboratory. No wonder, in the lab setting, actions look simply like events that *happen* to the actor. They are seen to be 'effects'

Slide It is because our actions are so irreducibly complex, the simple notion of 'a cause' - cerebral or otherwise - loses its grip on them so that they are not effects and even the more sophisticated notion of 'motive', understood as a

force external to, and certainly that of instinct, cannot easily be applied. What is the cause of your attending this talk? What *material* cause would you invoke? You may say: swathes of my entire past. But such a swathe – know-how, know-that, reasons, motives - is hardly a cause; and, if it were, it would be interesting to know who or what gathered it up so that it was able to act as a single cause. If it was *I*, then we are a long way from the notion of causation of my actions being something somehow outside of me. The idea of myself as a cause of my actions, or of my actions as an expression of myself, is close to the idea of freedom.

Let us dwell on this a bit more. While we concede that our past – and the sense of the future it informs - is deeply implicated in our actions, it is equally a mistake to think of ‘the past’ or a subset of past states of me as a mere cause of which we are passive effects. For a start, the past, or parts of it, is there as an explicit presence. Just think of the million components of know-that and know-how necessary for me to learn about Dr Libet’s experiments, decide to participate in them, and succeed in doing so. The temptation to see all this as a deposit of *effects* in my brain that then become causes can be resisted when we realize that they have to somehow be brought together to act upon me. The only way of synthesizing the disparate elements, so that they operate as occasions for ordinary actions, is through a sustained, forward-looking, explicit intention; in short, not through causes pushing from behind but through reasons pulling from in front. Reasons do not grow out of some putative biological substrate but are a forward-looking affirmation of, assertion of, expression of, myself.

The countless events which are subsumed in reasons cannot be generated – requisitioned, orchestrated - by ordinary causation or by processes of the kind that are described in neuroscience. It is wrong, for the same reason, to imagine that the orchestration of a multitude of movements, thoughts, rooted in knowledge and emotion, could be achieved by biological drives or motives that are themselves seen as quasi-material causes. Wishes, intentions, and other propositional attitudes are not simply caused nor simply causes. Like the actions that can be explained to some extent with reference to them, they are portions of a self-world that is more or less of a piece with other parts of the self and its world.

It is easy to see why committed determinists, including neuro-determinists, want to think of actions as *caused*: it prepares them to be reinserted into a causal chain extending backwards from a present material event to the Big Bang. But this is wrong. Yes, a journey to Birmingham to attend a meeting is a succession of movements; but it is more than that, which is why there is a difference between moving and travelling. Actions are not – and could not be - caused in the narrow, atomic, linear sense implied in the term ‘cause’ means. To see actions aright, we have to invoke the notion of an explicit purpose, which pulls us towards goals we have ourselves envisaged and articulated, and shapes the succession of action-components we undertake. This is the hidden nerve of association gluing together the myriad ‘sub-routines’ that make up components of actions, the countless elements that make up ordinary-sized actions (such as taking a train to Birmingham) and the innumerable actions that make up our lives, which we consciously and often effortfully lead rather than merely organically or material live or experience. To reduce reasons to, or to absorb them into, mere surface expressions of material or biological motors such as ‘motives’, ‘instincts’, or ‘drives’ (never mind unconscious ditto) is not only to misrepresent them but also to remove their explanatory force and to deprive complex but utterly ordinary actions of any kind of explanation. The action of ‘going to the Royal College of Physicians to make a case for improving epilepsy services’ as I have done on many occasions, could not be driven in this way because its goal and content are utterly singular and are rooted in my private, understood, recollected, past.

You might want to object that nest-building, such as birds engage in, likewise requires many thousands of moves that are not stereotyped. There is no a particular muscular signature, for example, corresponding to finding a wisp of straw to weave into the wall of the nest. Even so, there is a fundamental difference between this kind of complex instinctive behaviour and our everyday actions. Firstly, the overall action of nest-building *is* stereotyped at a very clear and simple level. The creation of a dwelling of a highly standardized form that has a clear function does not have to be articulated by the organism: nature takes care of that. Secondly, and connected with this, each of the elements is cued in by the previous element and the non-stereotyped components are clear instances of definite types. They do not require sustained intention informed by an explicit goal. *This* is the fundamental difference between plan-driven holiday-making and instinctive migration. The many components of the former, unlike those of the latter, are justified by, make sense with respect to, each other. The components are subordinated to explicit over-arching goals. Any voluntary action is a part of a nexus of behaviour that extends over swathes of what we might call ‘am-soil’ or ‘I-territory’ – in this respect quite unlike the events that comprise an epileptic fit or an animal’s programmed ‘courtship’ ritual.

Slide What have we established so far? Actions are *expressive* of a self operating in a world rather than a series of motor events arising in an organism, out of prior material event. Their organisation requires that they should be transilluminated by conscious intentions that make sense to the actor in the light of his/her past life and envisaged future. But will this satisfy all determinists?

Slide Am I justified in saying 'The buck starts here'? How can a material object in a material world be a point of origin of events? And even if my actions are an expression of myself, am I free if I did not cause or bring myself about? Let me deal with the first question first. **Slide** If you believe that you are identical with certain physical events in your brain, you will find it difficult to see where your freedom could be located; where there could be such a thing as the *initiation* of an action. Think about it: there is a sensory input, triggering nerve impulses which in turn trigger central nerve impulses that trigger motor activity or other outputs. Yes, there are many intermediate layers of activity between the input and the output, but they consist only of other nerve impulses and these are not qualitatively different from those more immediately related to inputs and outputs. Consider nerve impulses ascending through various tracts to the sensory cortex, being processed at higher and higher levels, and then, via various intermediate stations, activating the motor cortex, prompting outgoing impulses that descend to those fibres whose output causes muscles to contract. This sequence does not have a *beginning*, a point of *origin*, a point of *departure*, that would correspond to the initiation of an action. We have a loop of events passing through the nervous system, without an obvious point where anything could be *started*. After all, the circuitry of the brain is causally connected with its immediate surroundings and these are in turn simply part of a boundless causal nexus extending backwards in time to the beginning of the universe. The inescapable consequence of seeing ourselves identified with a material object – the brain – or a subset of activity within some part of it such as the neocortex - must be to conclude that we are wired into the material world, subject to the same laws which hold sway over it. Multiplying the number of intermediate steps between input and output, with so-called feedback and other loops, does not alter this: the causal chain gets longer and more complicated, that's all. Our seemingly voluntary actions still boil down to the inflection neural tissue gives to the flow of energy through a certain locality in the material world.

So a defence of freedom requires that deny that we are identical with activity in the brain, notwithstanding that the latter is a necessary condition for every aspect of consciousness from the slightest tingle of sensation to the most exquisitely constructed sense of self. The grounds for doing so is the theme for an hours' talk in itself but let me indicate roughly where the case for not identifying ourselves with neural activity lies. At its heart is the intentionality or aboutness of perception. **Slide** This is shown on the slide, with the upper line indicating the inward causal relationship between events in the perceived object that are essential to its being perceived and the lower line the intentionality that makes the perception (and, according to the neural theory of perception) the nerve impulse be *about* its object. **Slide** The inward causal chain explains how the light gets into my brain but not how this results in a gaze that looks out. This is a crucial step in the case against a neural or material account of consciousness but that is not my concern at present. **Slide** The bounce-back that we see on the slide, as we move from light getting in to a gaze looking out, and which marks the point at which perceptions are received is not merely the basis for the demarcation between input and output, it is also the first step towards the distinction between subject and object and the self and its world. It is out of this that there arises the possibility of the self as a 'centre' of a material world that has no centre and as a point of origin, a place where a buck could start. **Slide** This is too complex a theme to enter into but suffice it to say that intentionality, which tears the hitherto seamless fabric of a causally closed material world, is the seed out of which grows first person being and through, a trillion cognitive handshakes, creates a human world which in turn adds up to the virtual outside from which it is possible to act on a causally closed world.

That is far too quick and too much for anyone to take on trust and I hope we shall discuss this in question time. That is what I mean when I say that my actions are free in the sense of being *expressive* of myself. They belong to a field of action that is unique to myself; make sense only with respect to a frame of reference, a 'present past', a 'present future'; are rooted in 'am-soil', 'I-territory' related in turn to 'we-soil', 'we-territory'. In short they truly are manifestations of self-assertion or self-expression.

Which brings me to a subsidiary point, relating to the temporal depth of the self. Reinserting our actions back into the chains of causes and effects that make up the material world by breaking them up into small components not only reduces them to mere movements but also imprisons them in the present tense. It denies their *temporal depth* which has at least two dimensions.

The most obvious is the forward and backward connectedness of the components necessary to make sense of them and hence to make sense of the fact that they have occurred. For example, the complex movements involved in

locking and bolting my front door are intelligible only in relation to my having come out of my house and my intention to leave the house empty (so that it is vulnerable and needs protecting) while I am in London. And there is the additional dimension, which comes from the past self and the envisaged future self from which the trip to London draws its meaning and motivation. These two temporal dimensions are a particularly sophisticated elaboration of the intentionality or aboutness that characterizes my consciousness; indeed it is the orchestration by this aboutness that links the components of my action in (to borrow a phrase from Roger Scruton talking about music¹⁰) a 'virtual causality', a connectedness that is in another sphere from that of the material interactions of (say) my body and the pavement necessary for me to be able to propel myself to the railway station. It is easy to overlook the hinterland of self, the massive, tangled back-story, behind behaviour if we focus on individual actions lifted out of their context.

Slide It is easy to overlook this and we don't have to be sealed up in a lab to make this mistake. Consider catching a ball. The more brilliant the catch, the less it seems voluntary. We seem to have done it without thinking about, without *deciding* to do it. Indeed, when you consider what catching a ball involves, it seems impossible performing it as a *voluntary* act. You have to fling yourself across empty space in such a way as your outstretched hand intercepts the ball. The hand has to be sufficiently open at the time of contact as to admit the entry of the ball but not so wide open that the ball escapes. The fingers then have to close rapidly round the ball. You also have to allow a certain amount of compliance so that the ball does not at once bounce out of the hand before you have managed to trap it between your fingers. There are many other variables that have to be fixed, none of which you could deliberately control. So surely *you* did not catch the ball, *your body did* and you were just a fortunate bystander who took the credit.

No-one really thinks this – and for good reasons. First of all, in order to catch the ball, you had to participate in a game of cricket. This requires that you should have (voluntarily) turned up to a particular place on a particular day, that you understood and assented to the rules of cricket, and that you understood the role of the fielder, in particular that of the slip fielder. More importantly, in order to make the catch, you would have had to practise. This means hours spent in the nets, preparing yourself for this moment that would bring such glory upon you. This requires a sense of the future in which the present would influence what has become its past. You would have so to order your affairs that you would be able to go to the nets at the booked time – negotiating the traffic, making sure your day was clear so you could take up your booked slot, and so on. You would listen hard to your coaches' advice and do your best to translate it into action.

In other words, behind this quasi-involuntary action there would be a huge and complex hinterland of actions that could not have taken place without your deliberate intent. Over the months, you have carried out a vast number of voluntary actions so that you might be able when required to perform an action that you could not carry out entirely voluntarily. Many of these preparatory actions have taken the form of positioning yourself to have experience and acquire knowledge, deploying many intermediate steps in doing so. You toil in the nets in October with your inner eye on the May afternoon when October will have become a past from which you now benefit. And this is how it is with much of our life, which consists of acting on ourselves in order to change ourselves – from going to a pub to have a drink to cheer oneself up to paying good money to improve our chances of cutting a figure in Paris by polishing up our French.

You may think this is so obvious that it hardly needs to be spelled out but it is important not to underestimate the extent to which neuromaniacs overlook the obvious. Consider a recent study by Jan Scholz and his colleagues.¹² The researchers found that people who learned to juggle over a period of six weeks had clear changes in the white matter of a part of the cerebral cortex (the intraparietal sulcus) that is associated with visuo-motor skills. One of the authors, Heidi Johansen-Berg concluded from this that 'it's possible for *the brain* to condition its own wiring system to operate more efficiently' [italics mine]. In fact, it is not the *brain* that is doing this but the *participants* who enrolled in the experiment, committed themselves to training to juggle. They would have to remember to go upstairs to practise every day, to look after the juggling balls, set time aside for this purpose; in short, to engage in a set of actions of immense complexity that would not have been sewn together except by an individual who had a sustained and conscious intention to conform to the protocol of the experiment. The experiment, in short, provides clear evidence that it is true that *we train our brains rather our brains train themselves*. The trainer, in short, is not the brain but the person. The whole enterprise involved a large number of individuals, including the person who set up the experiment, an understanding spouse, children who kept quiet and played nicely while mummy was practising her juggling, and so on. The experiment by Scholz et al demonstrates how *persons* (not brains) increase their own agency by deliberate training – something, as we have already noted, no animal does. To ascribe it to the brain is a perverse spin. The acquisition of the skill was not a brain-directed plasticity of the brain but a person-directed plasticity of a person – interacting with the society which is the arena of the self.

While it is not entirely misleading to describe the acquisition of a skill (or the bodily basis of a skill) in neurological terms or to talk about neuroplasticity, we need to be reminded that neuroplasticity is often person-driven and that the person that does the driving cannot be understood without invoking the collective and individual transcendence that is the world and the self. We should not be so impressed by *neuroplasticity* that we forget bodily plasticity, plasticity of consciousness (including that increased confidence in our abilities which can be self-fulfilling), plasticity of the self, and, yes, plasticity of the world – as when I decide that others should work with me in a different way to ensure that one or other of us holds that so-important catch. It is a mistake to try to stuff all that back into the brain and see it in terms of changes in synaptic connexions at the microscopic level or alterations in cortical maps at the comparatively macroscopic level.

Slide The appeal to intentionality and first-person being and the self-shaping self as the *source* of free actions may still not impress some determinists.

But does this really prove that it is possible for events that are actions to have a place of origin – the self – without their being either mere effects of a material cause or mysteriously uncaused causes? What, anyway, is the standing of the self to which I refer? Isn't this also stitched into the world? Isn't it merely a set of effects of events that have impinged on it or its body?

One of the clearest and most succinct recent cases for hard-line determinism most relevant here comes from Galen Strawson who treats the self as an item in the causal net and has argued that, as we are not self-caused, freedom is impossible and moral responsibility is consequently groundless. **Slide**

Strawson's argument is very simple:

- a) Nothing can be the cause of itself.
- b) In order to be truly morally responsible for one's actions, one would have to be the cause of one's self.
- c) Therefore nothing (and hence no-one) can be truly morally responsible.

In order to be able to perform an act for which we are truly morally responsible, we would, it seems, have to be self-determining and this is impossible because the notion of true self-determination runs into an infinite regress. Supposing I choose my actions on the basis of certain principles. Where did those principles come from? If they were foisted on me, then I am not free. But suppose these principles were not foisted on me but I had a second set of principles to justify my choosing the first set. Then I would require a third set of principles to justify choosing that second set. And so on.

Strawson's argument is useful because it makes clear the assumptions behind determinism and, incidentally, reduces them to absurdity. In order to escape being determined, it seems, I have to have brought myself into being – a trick that of course only God can pull off. In order to be responsible for anything I *do*, I have to be responsible for everything that I *am*, including my very existence. Given that I cannot pre-exist my own existence, in order to be able to bring it about, this is a requirement that cannot be met.

It is clear that I cannot be a cause of myself if cause is understood as something like a material event. But I hope it will be clear from what I have just said that the acting self is quite different from the effect or material causes subsequently becomes a material cause itself in the way that a pebble is the effect of geological events and itself is the source of other events, as when it rolls down the hill. Our actions, however, though they operate on the material world do not originate in it: they arise, as already noted, from a different substratum – from the soil of the self-world. This is ultimately grounded in the Existential Intuition of first-person being – the sense 'That I am this' – where 'this' in the first instance is our own body though it develops beyond this. We appropriate our own bodies and by this means we are inserted in the world that exists *for us*. Our human world of pooled transcendence creates a theatre for our actions. Given that my actions have grown out of all those items, events and processes that I have appropriated – beginning with my body – in the service of my evolving and increasingly self-conscious, other-conscious and world-conscious ends, they have emerged from a soil that I *more or less* am: less, or hardly at all, as an infant; much more as an adult.

This is sufficient *causa sui* for me to be justly held responsible for my actions. To put this slightly differently: The first person is self-appropriating and its actions are ultimately rooted in the unfolding of the primary act of self-appropriation – the Existential Intuition that makes 'is' into 'am' and sets the 'I' off from the world which is the theatre and substrate of its led life. No-one at any rate can gainsay my intuition that my body is me and its actions mine. This is the version of *causa sui* that should answer anything meaningful in Strawson's demand. If we were to interpret this demand that, in order to be free, there should be *nothing* 'given' about ourselves, then freedom would be reserved for entities that were nothing and had nothing to be free about. Without a starter pack of the given, there would be nothing to be free, nothing for it to be free for, and nothing for it to be free from. A rather empty account of freedom, one would have thought,

though this is the implication of Strawson's position: that being anything that one has not manufactured oneself is a constraint on freedom. In fact, freedom does not require that we should be free of the given – that, for example, I shouldn't have a particular body that began at a particular time – but that we should take the given and run with it. Of course we cannot be free if we all contents of ourselves as alien bodies or external forces

Our self is neither a thing (like a pebble which, to an observer, has causal 'inputs' and effects as 'outputs') nor is it a mere succession or shower of material events. What I *do* makes sense with respect to a narrative that is my actively led life. Now it may be argued that we merely narrate what was going to happen anyway: we are deluded in the belief that we are free; as neurodeterminists might express it, the brain calls the shots and we retrospectively claim them as our own. That argument is in part dealt with by looking, as we have done already, at the nature of actions. But a lingering suspicion may remain that we can't be free unless we can somehow break the laws of nature. Only in this way could we deflect the course that events were going to take anyway. We need to address this seemingly insuperable objection to the claim that we are free, the materialist bedrock upon which biological determinism ultimately rests. And to do so, I want to borrow an idea from John Stuart Mill, one that he put forward in a paper published posthumously.

Mill was greatly exercised trying to reconcile his materialism with his passion for liberty. How can there be free agents, when we are material parts of a material world and subject to the laws of nature? He agreed that, yes, we have to obey the laws of nature – indeed, there is no choice. But we should appreciate that at any given juncture, there is more than one law of nature operating. By aligning ourselves with one law, we can use nature to achieve ends not envisaged in nature:

Slide Though we cannot emancipate ourselves from the laws of nature as a whole, we can escape from any particular law of nature if we are able to withdraw ourselves from the circumstances in which it acts. Though we can do nothing except through laws of nature, we can use one law to counteract another.

Slide We *utilize* the laws of nature by aligning ourselves with the one that leads to our goal and we do so from a virtual outside-of-nature that is the world opened up by intentionality. This virtual outside-of-nature is the realm I alluded to earlier: a human world whose seed is intentionality created out of a trillion cognitive handshakes sharing and elaborating intentionality. This public sphere, which is a dense network of signs of meant meanings, and a 'technosphere', in which we live and have our being beyond the material of our body, is where we first elucidate the laws of nature and get them to work on our behalf. It is where we use our pooled outside and our pooled strength to operate on the material world. This outside gives a place in which to step back. The stepping back is a huge collective stepping back into a space collectively created.

Slide Let me illustrate Mill's idea with a trivial example: going to a park in order to enjoy slithering down a slide, though it is diffusely present in the volition that pervades all our waking moments. The descent is courtesy of the laws of gravity but positioning ourselves to enjoy the descent is something else. Mummy has to agree and find the time. The trip to the park has to be organized, other things have to be fitted around it, there is a journey to the park, to the playground, and thence to the slide, guided by know-how and know-that, and there is an ascent to the top of the slide. The slide itself has been erected in order explicitly to utilize the laws of motion: it is a standing possibility of the joy of safely succumbing to the gravitational field – by appointment. This trivial example illustrates how our ways of acting involve knowledge, as well as artefacts (which of course operate within the laws of nature) so that we can subordinate them to our own ends and can, as Mill said, quoting **Slide** Francis Bacon, 'obey nature so as to command her'. Our actions are not uncaused miracles: they go with the grain of causation. But we are able to step back into the great extra-natural space that is the human world and from there use material causes as handles or levers on the material world.

Slide The most obvious manifestation of this is our exploitation of the laws of nature in science-based technology, a supreme expression of accumulated knowledge that is the property of the great community of minds. Technology is possible because we approach nature from that outside whose seed is intentionality. This outside is built up as an expanding Space of Possibility, a first-person plural reality, constructed through the joined endeavours of the human race, and expanded since the first hominids first awoke to their own existence. Such conscious exploitation of the laws of nature lies beyond description in terms of material causes and material effects: it cannot be described in terms of biological tropisms or instincts or drives as proxy for intermediate material causes.

Slide Let me draw to a close by looking at the claim that we really are able to act freely and how I have defended this claim by looking at how our actions meet the criteria for being free. When we think about the characteristics of a free act, three things seem to me to be paramount. First, the action should be *expressive* of what I am. This requirement is met by the fact that my actions are rooted in great swathes of myself – the 'am-soil' of which I spoke earlier. Secondly, I

should seem to be the *initiator* or source of my action. This is evident in the example of the slide or catching the ball. I carry out all the preparatory action to make a certain event happen, even one that ends up in helplessness (going down the slide) or automaticity (catching the ball). The counter-causal nature of intentionality, which lays the seed for our distance from the world, for our sense of self and our freedom, which shared or joined, is the basis for the human world offset from nature and it is this that makes us a point of origin, making each of us individually the centre of a centreless universe, a place where 'the buck can start'. Finally, my actions should *deflect* the course of events rather than merely conform to what was anyway going to happen. What evidence is there for such deflection?

Slide Anyone who doubts that we can individually deflect the course of events should consider what we have achieved in building up a human world so extensive as virtually at times to conceal the natural one. **Slide** As was said of Christopher Wren '*Si monumentum requiris, circumspice*': if you seek his monument, look about you. The artifacts of cities which cover the surface of the earth with man-made objects, the human institutions to which we relate for so much of our lives, and the extra-natural social facts and preoccupations that fill our waking hours, to which there is nothing corresponding in nature – these are eloquent testimony to how, collectively at least, we deflect the course of events and operate within a space outside of the material world construed according to the laws of physics. **Slide** From pointing, through artifacts and spoken, and ultimately written, language, we get ever greater purchase on the natural world from an ever greater outside built up by thousands of generations comprised each at first of thousands, then of millions and ultimately of billions, of people.

This should be enough to satisfy everyone that we are capable of truly free actions. There will still be some who are dogmatically opposed to the idea of our being free because it doesn't fit with what they believe to be the scientific world picture. To the latter, we offer this question: If freedom really is an illusion where on earth did the illusion come from, particularly such a tenacious one? **Slide** As Dr. Johnson observed 'All theory is against freedom of the will; all experience is for it'. Perhaps not all theory – just some theories of some philosophers – and clearly not all experience – just most of our everyday experience. Even so, like Dr. Johnson, you might be inclined to agree with his tetchy assertion that 'we know our will is free and *there's* an end on't'. **Slide** Why, after all, if freedom is an illusion, and such a stubborn one, and material causation reigns unchecked and undeviated, should one section of the infinite causal nexus of the universe decide, apparently without any foundation, that it is itself a point of origin of certain events - actions - that are not simply part of an endless chain of causes whose ancestry ultimately lies in the big bang? It seems an odd idea for a causal net, or a bit of it, to entertain.

There is (inevitably) a Darwinitic response. Anything is possible (even, as we have seen, consciousness itself) if it is of adaptive value. The reason Dr. Johnson can't get rid of the idea that he and his fellow humans are free, neurodeterminists argue, is because this will be good for his, and our, morale. The sense that I am the source of my actions gives me an enhanced potency and also, by making me feel responsible for certain events that I deem to be my actions, makes me ethically more biddable. I can, for example, feel shame. As Rita Carter says

Slide The illusion of free will is deeply ingrained precisely because it prevents us from falling into a suicidally fatalistic state of mind- it is one of the brain's most powerful aids to survival. Like many of our survival mechanisms, however, it no longer works to our benefit. By creating the illusion that there is a self-determining 'I' in each of us, it causes us to punish those who appear to behave badly, even when punishment clearly has no practical benefit.

Slide This is an interesting claim because it suggests that our belief that we are free *can* (after all) alter what happens in the world – initially, so far as we are concerned, for the better because it helps us to survive. In short, the illusion of free will *does* deflect the course of events – and hence is self-fulfilling. It is *not* an illusion. For if we really cannot deflect the course of predetermined events, and then the idea that we are free cannot change anything, any more than the idea that we are not free can change it.

Well, I have tried your patience too long. I hope you are persuaded that it is possible to accept that we are capable of free actions, in the sense of events that are expressive of us, originate with us, and deflect the course of things; that they do not require us to break the laws of material nature; and that neuroscience adds nothing to the flawed case for believing that free will is an illusion. Neurodeterminism works within the same assumptions as determinism period but it usefully highlights the latter.