

The papers of Strawson, Steward, and Duff thus represent three different views on the compatibility of physical determinism and our responsibility for our actions. Steward is an 'incompatibilist', holding that physical determinism would rule out moral responsibility. Duff is a 'compatibilist' (with respect to the criminal law), claiming that we could still rationally hold people legally responsible for their actions even if physical determinism were true, while Strawson claims that we could never be (ultimately) responsible for our actions, whether or not determinism is true.

# 1

## Does brain science change our view of free will?

PATRICK HAGGARD<sup>1</sup>

---

### 1. Introduction

Free will represents a frontier for Neuroscience. A neuroscientific account of free will would require a reductive explanation of something that we recognize as an important, even constitutive feature of ourselves as persons. In recent years, the possibility of a neuroscientific account of free will has attracted growing public attention. It now rivals classic neurophilosophical topics such as visual awareness in terms of public visibility. This may be partly because free will has a more direct importance to society than visual awareness: an idea of free will underlies the concept of responsibility for action, on which our legal and social systems depend. Our growing understanding of the brain mechanisms of human agency has brought two ways of thinking about persons into apparent conflict. The neuroscientific view focuses primarily on brain mechanisms: behaviour, decision and individual consciousness are all consequences of these mechanisms. On the other hand, the everyday, folk-psychological concept of free will is based on our first-person experience of choosing and acting, plus a strong component of social-cultural norms. These two views conflict in a number of ways. For example, if neuroscience were to invalidate our concept of free will, what would be the effects on our societies and ways of life? Less dramatically, if neuroscience were to show that free will resulted from a specific brain mechanism, how would society respond to actions of an individual with damage to that mechanism?

<sup>1</sup> Preparation of this chapter was supported by a Research Fellowship from the Leverhulme Trust, and by an ESF/ESRC research grant on 'Intentional inhibition'.

To begin, I want to declare two specific perspectives that this chapter will take. First, I will consider free will from the perspective of the human *motor* system. This view is justified because the problem of free will is a problem about the causes of our motor *actions*. For example, we might want to know what led to a person pulling a trigger, choosing to utter a particular word, or even eating a particular apple. The approach from motor neuroscience begins with the fact of the motor action, and then traces back to identify the particular stages of information processing in the brain that eventually produced it. The argument progresses backwards from actions towards intentions, the reverse of the natural process that leads forwards from intentions towards actions. The reason for this direction is simply methodological. Neuroscience needs to begin with the measurable facts of behaviours. Second, I will not engage with the philosophical debate over whether free will is compatible with determinism or not. Neuroscience is fundamentally deterministic, in its methods, its assumptions and its outlooks. Therefore, the perspective from neuroscience requires that free will must either be compatible with determinism, or we do not have free will.

This chapter explores the interaction between neuroscience and free will in the following way. First, I will consider how freely willed actions should be defined. Second I will outline our current understanding of brain mechanisms preceding action, showing in what respects these mechanisms meet the philosophical criteria for freely willed action, and in what respects they do not. Finally, I will end by concluding that the philosophical criteria themselves are based on two underlying psychological facts: human action involves complex mappings between environmental stimuli and goal-directed responses, and human action is associated with a range of quasi-perceptual experiences, classically called 'motor attention'. These facts lie at the heart of our concept of conscious free will, and are directly related to the recent evolutionary development of the brain's frontal lobes.

## 2. Criteria for freely willed action

These criteria for free will come from philosophy and folk psychology. We may therefore ask how these criteria fit with our understanding of the brain. If there are any freely willed actions, how are they generated in the brain?

Two characteristic features seem to define freely willed actions. The first is the well-known 'Could have done otherwise' feature. In philosophy, this is often used as a test of compatibility with determinism: given the pre-existing state of the universe, could the action have occurred differently, or not occurred at all? Movement neuroscience understands 'could have done otherwise' in terms of a process of action selection, that arbitrates between competing action alternatives, and selects the one that eventually occurs. So, for example, if we have evidence from brain measurements that two different action alternatives were developed in parallel, up to a late stage just before action execution, and if we can identify the particular brain process that led to one action alternative being selected and the other deselected, we may say that the person, or animal, could have done otherwise. The mechanism of selection itself may in fact be deterministic: that does not concern us here. The important point from a neuroscientific perspective is that the system supports a range of possible actions, rather than just one, and contains a mechanism that selects definite actions, and eliminates alternative possibilities. Whether this is also a reasonable philosophical interpretation of 'could have done otherwise' is less clear. Much depends on the meaning of the word 'could'. The philosophical interpretation often carries the idea of whether an alternative possibility was open or not, given laws of nature and initial conditions. The neuroscientific interpretation emphasizes only that alternative possibilities were *represented* in the mind/brain, and does not comment on whether those possibilities could have actually happened.

The second characteristic of freely willed actions is more elusive. Our concept of free will implies a conscious subject, who initiates the action, and is therefore responsible for it. That is, there must be an 'I' who could have done otherwise. Moreover, the agent must be conscious that they are making the action. We will therefore refer to this as the *subjective agency*. Therefore, concepts of free will do not apply when this subjectivity is absent, for example to a person who acts in their sleep, or to an involuntary movement, such as a simple reflex. Neuroscience treats the 'I' as being synonymous with an individual's brain. Therefore, it approaches the awareness of action by distinguishing between those action-related processes in the brain that typically generate conscious experiences, and those that do not. The neuroscientific view of *conscious* free will accordingly pays particular attention to the relation between brain activity in the motor system and the experience of intending and performing an action.

### 3. Methodology: how can volition be studied experimentally?

Before examining the contributions of brain science to our understanding of volition, we must first ask about methodology: how can volition be studied experimentally?


Voluntary action often appears to vanish in the controlled conditions of the experimental laboratory. In particular, good experiments require good specification of the stimulus, and clear instructions, which both seem incompatible with free will. Thus, an experiment with the instruction 'have free will, now!' seems absurd. The concept of free will implies that the participant themselves, rather than the experimental situation, brings the triggers for action. This impasse persuades many that volition cannot be studied scientifically. Indeed the concept largely disappeared from psychology during the behaviourist and cognitivist periods (Skinner 1953).

Experimentalists have generally escaped from this impasse by comparing 'fixed' actions, where an external stimulus specifies which of a range of possible actions a person must make, with 'free' actions, where the participant must generate this information for themselves. For example, Deiber *et al.* (1991) compared a condition in which participants were instructed to move a joystick repetitively in a single direction, with a condition in which the participants themselves selected the direction. These conditions both involved a physical movement, but only the second condition involved free selection from among action alternatives. Thus, comparing the two conditions should reveal the brain activity associated with voluntary action selection.

This 'fixed vs. free action' approach provides an operational definition of volition, based on externally specified vs. internally generated information for action. It has come to dominate most neuroscientific studies. Brass and Haggard (2008) argued that decisions about whether to perform an action, what action to perform, and when to perform it, could all be understood within this framework. This operational definition does appear to satisfy the 'could have done otherwise' criterion for free will, as described above. However, the subjective agency component is still missing. Contrasting forced vs. free actions tells us nothing about the subjective aspect of action selection, nor about why and how one particular action was chosen over another. The participant clearly generates action information internally, but we do not know how. One speculation is that people try to choose randomly, in as unfixed a way as possible.

### 4. Brain pathways underlying voluntary action

The idea of voluntary action is intimately tied to the control of muscular movement. Wittgenstein asks: 'what is left over if I subtract the fact that my arm goes up from the fact that I raise my arm' (Wittgenstein 1953). Therefore, the logical way to explore the brain pathways underlying voluntary action is in the backwards direction, beginning with the brain area that sends the final motor command to the muscles: the primary motor cortex. The primary motor cortex forms a strip of tissue immediately in front of the central sulcus in each cerebral hemisphere (Figure 1.1).



**Figure 1.1** Cortical structures involved in generation of voluntary action.

M1 – primary motor cortex; SMA – supplementary motor area; preSMA – presupplementary motor area.

Source: P. Haggard (2008) 'Human volition: towards a neuroscience of will', *Nature Reviews Neuroscience*, 9: 934–46.

## 5. Primary motor cortex: the final common path

The primary motor cortex represents the major motor output area of the cerebrum: the neurons here communicate directly with the lower motoneurons in the spinal cord, and so are only one synapse away from the muscle itself. Interestingly, in recent years it has become possible to activate this area directly from outside the brain using transcranial magnetic stimulation (TMS). The brief magnetic field induced by the machine (Figure 1.2) passes through the skull, and causes neurons in the primary motor cortex to fire. About 20ms later, depending on the distance that the nervous impulse has to travel between brain and muscle, there is a small involuntary twitch of the hand, which can easily be measured using electrodes placed over relevant hand muscles. Interestingly, the experience of TMS is absolutely different from the experience of voluntary action: the participant quite distinctly feels that they are *being moved*, like a marionette is moved by pulling a string. This result immediately suggests that the primary motor cortex cannot form part of the brain mechanism underlying free will, because it does not satisfy our criterion of subjective agency. Indeed, if stronger TMS pulses are used to delay a voluntary action, the participant does not perceive that their action has been delayed. This suggests that the experience of voluntary action has been formed prior to an earlier stage (Haggard and Magno 1999).

The brain mechanisms of free will must therefore lie upstream, in the areas that initiate action, and not in the final output stage that dispatches motor commands to the muscles. Sherrington developed the classic physiological concept of a 'final common path' to action (Sherrington 1906). This states that, while there are many reasons and origins for our actions, they must ultimately converge on to the specific centres that cause our muscles to contract. The primary motor cortex is part of the final common path that leads to muscle. But the discussion above shows that our criteria of 'could have done otherwise' and subjective agency apply at the beginning of the final common path, closer to the reasons for action, and not at the end.

## 6. The supplementary motor areas underlying action selection

Therefore, pursuing the motor chain backwards from the primary motor cortex should help to identify brain areas underlying voluntary action.




Figure 1.2 Transcranial magnetic stimulation over the motor cortex (upper panel) can be used to cause involuntary movements of the hand (lower panel). Note the electrodes placed over the muscles that primarily move the index finger. The deflection of the trace on the computer screen shows the twitch caused by the stimulation.

Here, there has been strong interest in areas immediately in front of the motor cortex, notably the supplementary motor area (often termed SMAproper) and the presupplementary motor area (preSMA). These areas will therefore be discussed in detail, because they seem to come close to our criterion.

To understand what a brain area does, neuroscientists often begin by noting how behaviour changes when the brain area is damaged. They then conclude that the function of the brain area corresponds to whatever the person or animal cannot do following localized damage. There is convincing evidence that the preSMA is involved in the selection of alternative actions, and that it achieves this selection through competition between action possibilities. A dramatic instance of this comes from Anarchic Hand Syndrome (AHS), which is occasionally seen after damage to the SMA in one hemisphere, particularly when the callosal fibres that carry signals between the two hemispheres are also damaged (Della Sala *et al.* 1991). In AHS, the patient performs unintended, goal-directed actions with the affected hand, as if the hand is driven by the immediate stimulus environment, rather than by the patient's personal will. Della Sala *et al.* give the compelling example of a patient saying that they will wait before taking a drink, because the beverage is probably still too hot to drink, while the affected hand compulsively moves towards the cup, to grasp it and bring it to the mouth. Patients with small lesions in this area, but without full AHS syndrome, are unusually responsive to subliminal priming by visual stimuli, as if their action is captured by current sensory information (Sumner *et al.* 2007).

The broad picture that emerges from this literature is of environmental stimuli that continually elicit competing responses. A major part of volition involves the inhibitory process of deselecting and suppressing those stimulus-triggered responses when they are inappropriate, or not useful for our current goals. Volition involves exerting internal motivations against external triggers for action.

## 7. Readiness potentials and action initiation

This data can explain the negative side of volition (why we do not constantly react), but says less about the positive side (why we sometimes perform actions that relate to our goals and wishes). Prior to voluntary actions, a characteristic pattern of electrical activity occurs in the SMA and

preSMA. This can be recorded by electrodes on the frontal scalp as a 'readiness-potential', or gradual ramp-like increasing negative voltage, beginning 1s or more before action, and ending abruptly around the time of action itself (Kornhuber and Deecke 1965). Associated with growing premotor activity in frontal motor regions is a decrease in the spontaneous oscillatory activity of the same regions (Neuper *et al.* 2006). The conventional recording method, with electrodes placed on the scalp, does not have enough spatial resolution to identify precisely the underlying brain regions. However, intracortical recordings in neurosurgical patients with electrodes implanted in the brain confirm that the preSMA and SMAproper are involved (Ikeda *et al.* 1999).

In fact, recent neuroscientific data suggest that the positive process of initiating action is intimately linked with the process of action selection. Our 'could have done otherwise' criterion focuses on the situation where the person decides for themselves on the details of action. Interestingly, when a person is asked to freely choose between an action of the left or right hand, the later stages, but not the earlier stages, of the readiness potential show a greater voltage over the right and left hemispheres of the brain, respectively, reflecting the fact that muscular control of each side of the body depends on the motor areas in the contralateral hemisphere. Thus, the readiness potential is divided into an early non-lateralized phase, where the participant might presumably make either a left or right hand action, and a later, lateralized phase where these details of action are clearly more fixed.

Further evidence of a link to action selection comes from a recent study of premovement reduction in neural oscillations that accompanies the readiness potential (Tzagarakis *et al.* 2010). When participants were warned in advance that they would shortly have to move in one of several possible directions, the oscillatory power decreased as the number of possible movements decreased. For example, when participants knew they need select from only two possible movement directions, oscillatory power was more strongly reduced than when they were instructed that they would select from three.

The readiness potential, and associated reductions in neural oscillations, are important measures of the processes that *initiate* voluntary action. These processes can be divided into an early phase, which does not fully specify the details of which action will be performed, and a later phase, which contains more motor detail. The neural events that precede action can be seen either as an energization process that triggers action, or as a

progressive reduction of uncertainty that selects a specific intended action, and deselects alternatives. While we experience our voluntary action as a positive process aiming towards a goal, neuroscientific data show that initiation of voluntary action is associated with a neural process of action selection. Indeed, to make a voluntary action may simply be to reduce the size of what has been called the 'response space', to converge on a single action. In terms of the 'could have done otherwise' criterion, there is strong evidence for a dynamic process that gradually evolves from a situation where multiple actions are possible, to a point of no return, where a single action enters the final common path in the primary motor cortex. One might say that in voluntary action, we move from a phase where we clearly could do otherwise, to a point where we are prisoners of our current motor command. We may be free before we act, but once we act, counterfactuals become irrelevant. Responsibility for action requires that actions, in the end, are objective, physical facts about the world.


## 8. Brain processes underlying the experience of volition

We now turn to the second criterion, of subjective agency. The phenomenology of volition is hard to study experimentally. We know when we make voluntary actions, and we certainly have no difficulty discriminating our voluntary actions from physically identical passive movements, as in Wittgenstein's example of raising one's hand. However, the feeling of voluntary action is often thin and elusive. I want to walk to the station, and I decide to do so, but I do not have an intense experience of volition with each step that takes me there.

The best-known approach to the conscious experience of willing is the chronometric approach, in which people are asked to report the moment in time when they experience the intention to initiate an action. Clearly, this method is restricted to the immediate, decisional aspect of volition that precedes action, and cannot capture general, longer-range aspects of volition, such as my intentions to be more generous to my friends, or to repaint the house. The classic experiment in this area is that of Benjamin Libet (Libet *et al.* 1983). Libet and colleagues asked participants to make simple voluntary actions at a time of their own choosing while watching a rotating clock hand. Participants were further asked to note the position of the clock hand at which they first felt the 'urge' to move. After the action itself, the clock stopped after a random delay, and the participant reported

the previously-noted clock time. Readiness potentials were measured from the scalp throughout.

The result of the experiment is straightforward (Figure 1.3). The readiness potential began up to 1s or more before movement onset. The experience of willing the action occurred on average 206ms before movement onset. Libet argued that conscious will could not therefore be the cause of our actions, since the preparatory brain activity that produces the action begins before the conscious state of willing, and causes cannot precede their effects. Incidentally, this method of measuring perceived timing has attracted considerable criticism (see for example the commentaries to Libet's 1985 target article in *Behavioural and Brain Sciences*). However, these arguments have been extensively reviewed elsewhere. Here we focus principally on the relation between brain activity and subjective agency, rather than on the specific temporal order in the original experiment. From a neuroscientific perspective, Libet's result is actually



**Figure 1.3** Schematic representation of the Libet experiment.

Source: P. Haggard (2005) 'Conscious intention and motor cognition', *Trends in Cognitive Sciences*, 9(6): 290-5.



unremarkable. Conscious experience is a product or aspect of some brain processes, and not an event independent of the brain. Although consciousness may be part of brain activity, consciousness cannot cause brain activity, nor can it cause actions.

An alternative reading of Libet's results might be that the readiness potential, once it reaches a particular threshold value, produces an experience of conscious intention-in-action, i.e., of being about to act. Haggard and Eimer (1999) used Mill's method of concomitant variation to investigate this possibility. They reasoned that, if the readiness potential is indeed the cause of the conscious experience of intention, actions where the participant happens, for whatever reason, to experience a very early intention, long before they act, should have earlier readiness potentials than trials where the participant happens to experience intention rather later, just prior to the action itself. This argument has an important advantage over temporal precedence arguments, namely that it is unaffected by systematic errors in subjective time perception.

In fact, Haggard and Eimer found no evidence for such a relation between readiness potentials and the perceived time of intention. Further, their experiment gave participants a free choice between making an action with the left or right hand. In contrast, Libet's original experiment involved only unimanual actions. As described above, when people choose between left- and right-hand actions, the readiness potential begins over both brain hemispheres, but then lateralizes to the hemisphere contralateral to the hand that will move. Haggard and Eimer found a significant relation between the onset time of the *lateralized* readiness potential and the time of conscious intention. When conscious intention occurred early, there was an early lateralization of the readiness potential. When conscious intention occurred late, just prior to the action itself, there was a late lateralization of the readiness potential. Taken together, these results indicate that the general readiness potential cannot be the cause of the conscious experience of intention reported in Libet-type experiments. Conversely, the lateralized readiness potential might potentially be the cause of conscious intention. This finding is interesting because of the role that it gives the 'could have done otherwise' criterion in subjective agency. Specifically, the experience of conscious intention-in-action occurs once the action has taken a concrete motor form, and therefore after the processes of action selection. The brain processes that choose from among a set of alternative actions seem to indicate the point at which conscious experience enters the chain of voluntary action.

## 9. Are experiences of intention just retrospective confabulations?

A common objection to these experiments is based on the comment that participants report the time intention only *after* they have acted. Therefore, they may reconstruct an experience of 'intention' as part of a narrative to explain the occurrence of the action. Intentions are really narrative reconstructions, or even confabulations, and play no part in the brain processes that lead to action. This view has been called 'The illusion of conscious will' (Wegner 2003). While the mind clearly does reconstruct such narratives, I want to highlight two pieces of evidence from neurophysiology and psychology, that the subjective experience of voluntary action cannot *only* be based on retrospective reconstruction.

The neurophysiological evidence comes from direct stimulation of the human brain. In some cases of drug-resistant epilepsy, electrode grids are placed directly on to the cortical surface for exploratory stimulation prior to neurosurgical treatment. In a typical exploratory session, the patient remains fully conscious while the neurosurgeon stimulates each electrode in the grid. In some patients, stimulation of the supplementary motor area produces a feeling of an urge to move. The urge refers to a specific body part, such as the left hand (Fried *et al.* 1991). Crucially, the patient reports this feeling even when they do not know where or if they are being stimulated, and while remaining quite still. Since the patient has not actually moved, it is unclear how their experience could be based on a reconstructive narrative. Rather, it appears that the neurosurgeon has given them, quite involuntarily, an experience which seems like a percept of volition. The will is clearly not free in this case, since the neurosurgeon controls it. However, this data provides powerful evidence that the normal experience of subjective agency in volition is a direct consequence of activity in the supplementary motor areas.

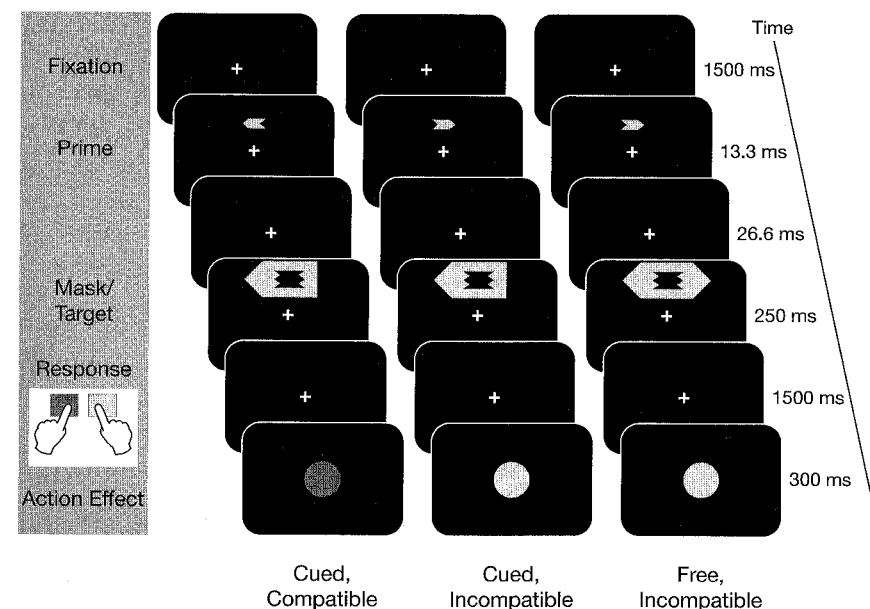
In fact, neurosurgical findings of an urge to move are rather widespread in the brain. They have also been reported in the parietal lobes (Desmurget *et al.* 2009), with which the frontal motor areas are strongly interconnected (see below). For our purposes, the precise location does not matter. The important point is that an external intervention within a specific brain mechanism can generate an experience which has at least some of the key features of volition.

The second class of evidence comes from psychological studies investigating the perceived feeling of control. Wenke and colleagues (including

the author) investigated whether the feeling of control over the effects of an action depends on the processes that initially select which action to make (Wenke *et al.* 2010). Since these action-selection processes necessarily occur before action itself, any such influence cannot be retrospective. Participants pressed a button with their left or right hand. In some trials, a large, clearly visible arrow showed which button to press. In other trials, a double-headed arrow appeared, meaning that they should freely choose for themselves which action to make. Without the participants' knowledge, the large arrows were preceded by invisibly brief presentations of smaller left- or right-pointing arrows. When participants were asked to freely choose between left and right keypresses, these 'subliminal priming' arrows were able to bias their 'free' choices, so that they were slightly more likely to choose the action compatible with the prime (e.g. left-priming arrow followed by a free choice to press the left key), rather than the incompatible action. When participants were explicitly instructed which key to press by the larger arrow, their reactions were slower and they were more likely to press the wrong key, if the large arrow had been preceded by an incompatible prime than by a compatible one. This shows that the priming arrows, though invisible, influenced the selection of the keypress actions, even when these were supposedly 'free'.

Each keypress was then followed by one of several possible colours appearing briefly on the screen. After many such trials, participants rank ordered all the different colours they had seen, in order of the level of control they had felt over the appearance of each colour. Crucially, we organized the colour patches so that subjects saw one set of colours when their action was compatible with the preceding subliminal prime, and another set of colours when their action was incompatible with the preceding prime.

The results showed that participants experienced a greater level of control over the colour when their action was compatible with a preceding subliminal prime than when it was incompatible. It is important to point out that the colour does not depend directly on the subliminal prime, but on whether the prime influences the subject's action selection processes or not. These results show that the initial, pre-movement process of selecting which action to perform makes an important contribution to the feeling of control over actions and their effects. Extrapolating to situations outside the laboratory, we can say that the sense of agency is not merely a post-hoc inference or confabulation to account for the actions that we find we have made. Rather, the sense of agency depends on the brain processes that lead



**Figure 1.4** The design of Wenke and colleagues' experiment. The leftward and rightward primes presented at the start of each trial are not consciously seen by the participant, though they bias how they respond to the subsequent instruction given by the large arrow targets. More interestingly, participants feel a stronger sense of controlling the subsequent colour patch caused by their action when they act in a way that is compatible with the unseen primes.

Source: D. Wenke, S.M. Fleming and P. Haggard (2010) 'Subliminal priming of actions influences sense of control over effects of action', *Cognition*, 115(1): 26–38.

us to make this particular action, rather than any other action, in the first place.

Wenke *et al.*'s experiment did not specifically ask about the experience of intention or volition, but only about the experience of control over effects of action. Interestingly, however, the results show that a pre-movement process of voluntary choice is highly relevant to the experience of control. This experiment also has interesting societal implications. The sense of control is strongly linked to the social concept of responsibility. For example, the law assumes that an individual has a reasonable awareness of the consequences of their action at the time that they perform the action. Wenke *et al.*'s experiment suggests that the sense of control depends, at least partly, on the way that the plan for action is selected and formulated



in the mind, prior to the actual physical execution of action. In that sense, to accept responsibility for one's actions is simply for there to be a particular functional relation linking the mental and neural states that occur prior to action, the physical movements of the body, and the subsequent external events caused by those movements.

## 10. Free will as action flexibility and complexity

At the start of this chapter, we suggested that brain processes of action selection may correspond to the 'could have done otherwise' criterion for free will. This view leads to the thought that the characteristic feature of freely willed actions may be appropriate engagement of the brain's action-selection system. Neuroscience makes a very general distinction between simple reflex actions, in which a single stimulus elicits a stereotypical response, and the more complex case, where a range of stimuli are mapped on to a range of responses, with a contextual rule defining the mapping between the two. Everyday life offers ready examples of such mappings: a red light means 'don't walk' if I am trying to cross the road, but a rather similar red light appears above the lift to tell me that the lift is here, and I can walk into it. This kind of rule-based mapping offers a kind of flexibility of action that simply does not exist for simple orienting responses, such as looking towards the source of an unexpected loud noise. Moreover, an action selection mechanism is required to map stimulus and context to action in the correct way. In humans, the rules mapping stimuli and contexts to appropriate responses can become extremely complex. The number of stimuli and contexts that are mapped is also extremely large. The key brain regions for implementing complex, flexible mappings are in the frontal and prefrontal cortex. The dorsolateral prefrontal cortex houses a form of working memory, that links the set of available stimuli to alternative actions, and possibly also selects from them. The medial frontal and prefrontal cortices seem to house a competitive selection from among a wider set of possible actions, stored in long-term memory.

Perhaps the key feature of 'free will' is simply the complexity and flexibility of the mappings in the brain that produce some of our actions. For freely willed actions, there often seems to be no obvious external cause or stimulus to explain why the agent performed one action rather than another. Dualist philosophies classically revert to some form of agent causation at this point. Nothing in the external world caused the action.

Rather, the agent themselves caused the action, even though they could have done otherwise. The neuroscientific view, in contrast, would suggest that everything, rather than nothing, is the cause of our freely willed actions. Such actions may simply be responses to a situation whose complexity seems to exhaust our attempt to identify a simple cause for each action.

## 11. Conclusion

Brain science will change our notion of free will in several ways. Most importantly, it will show that free will is not a special, transcendental faculty. Rather, it is a term given to the operations of a set of brain processes in the frontal lobes of humans, and possibly some other primates. More specifically, brain science will progress the philosophical analysis of free will in three distinct ways. First, it will show that the 'could have done otherwise' criterion of free will refers to the engagement of a neural memory buffer that stores possible alternative actions. Second, it will show that the conscious experience of initiating and controlling is a quasi-perceptual process resulting from neural activity in specific frontal and parietal brain areas. Finally, it will provide an alternative way of thinking about agent-causation. Conventional metaphysics of free will invokes an 'I' to consciously initiate willed actions. Instead, neuroscience emphasizes that 'freely willed' actions may simply have a more complex set of causes than simpler actions. Nevertheless, these causes are not special in any particular sense: they simply reflect the flexibility and complexity of our response to our environment.

## References

- Brass, M. and Haggard, P. (2008) The what, when, whether model of intentional action, *The Neuroscientist*, 14(4): 319–25.
- Deiber, M.P., Passingham, R.E., Colebatch, J.G., Friston, K.J., Nixon, P.D. and Frackowiak, R.S. (1991) Cortical areas and the selection of movement: a study with positron emission tomography, *Experimental Brain Research. Experimentelle Hirnforschung. Expérimentation Cérébrale*, 84(2): 393–402.
- Della Sala, S., Marchetti, C. and Spinnler, H. (1991) Right-sided anarchic (alien) hand: a longitudinal study, *Neuropsychologia*, 29(11): 1113–27.

- Desmurget, M., Reilly, K.T., Richard, N., Szathmari, A., Mottolese, C. and Sirigu, A. (2009) Movement intention after parietal cortex stimulation in humans, *Science* (New York), 324(5928): 811–13.
- Fried, I., Katz, A., McCarthy, G., Sass, K.J., Williamson, P., Spencer, S.S. and Spencer, D.D. (1991) Functional organization of human supplementary motor cortex studied by electrical stimulation, *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 11(11): 3656–66.
- Haggard, P. and Eimer, M. (1999) On the relation between brain potentials and the awareness of voluntary movements, *Experimental Brain Research. Experimentelle Hirnforschung. Expérimentation Cérébrale*, 126(1): 128–33.
- Haggard, P. and Magno, E. (1999) Localising awareness of action with transcranial magnetic stimulation, *Experimental Brain Research. Experimentelle Hirnforschung. Expérimentation Cérébrale*, 127(1): 102–7.
- Haggard, P. (2005) Conscious intention and motor cognition, *Trends in Cognitive Sciences*, 9(6): 290–5.
- Ikeda, A., Yazawa, S., Kunieda, T., Ohara, S., Terada, K., Mikuni, N. and Nagamine, T. (1999) Cognitive motor control in human pre-supplementary motor area studied by subdural recording of discrimination/selection-related potentials, *Brain: A Journal of Neurology*, 122(5): 915–31.
- Kornhuber, H.H. and Deecke, L. (1965) [Changes in the brain potential in voluntary movements and passive movements in man: readiness potential and reafferent potentials], *Pflügers Archiv Für Die Gesamte Physiologie Des Menschen Und Der Tiere*, 284: 1–17.
- Libet, B., Gleason, C.A., Wright, E.W. and Pearl, D.K. (1983) Time of conscious intention to act in relation to onset of cerebral activity (readiness-potential): the unconscious initiation of a freely voluntary act, *Brain: A Journal of Neurology*, 106(3): 623–42.
- Neuper, C., Wörtz, M. and Pfurtscheller, G. (2006) ERD/ERS patterns reflecting sensorimotor activation and deactivation, *Progress in Brain Research*, 159: 211–22.
- Sherrington, C.S. (1906) *The Integrative Action of the Nervous System*, New York: Charles Scribner's Sons.
- Skinner, B.F. (1953) *Science and Human Behavior*, New York: Macmillan.
- Sumner, P., Nachev, P., Morris, P., Peters, A.M., Jackson, S.R., Kennard, C. and Husain, M. (2007) Human medial frontal cortex mediates unconscious inhibition of voluntary action, *Neuron*, 54(5): 697–711.
- Tzagarakis, C., Ince, N.F., Leuthold, A.C. and Pellizzer, G. (2010) Beta-band activity during motor planning reflects response uncertainty, *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 30(34): 11270–7.
- Wegner, D.M. (2003) *The Illusion of Conscious Will* (First Edition), Cambridge, MA: MIT Press.
- Wenke, D., Fleming, S.M. and Haggard, P. (2010) Subliminal priming of actions influences sense of control over effects of action, *Cognition*, 115(1): 26–38.
- Wittgenstein, L. (1953) *Philosophical Investigations*, Oxford: Blackwell.

## 2 Libet and the case for free will scepticism

TIM BAYNE<sup>1</sup>

---

### 1. Introduction

Free will sceptics claim that we do not possess free will – or at least, that we do not possess nearly as much free will as we think we do. Some free will sceptics hold that the very notion of free will is incoherent, and that no being could possibly possess free will (Strawson, this volume). Others allow that the notion of free will is coherent, but hold that features of our cognitive architecture prevent us from possessing free will. My concern in this chapter is with views of the second kind. According to an increasingly influential line of thought, our common-sense commitment to the existence of free will is threatened in unique ways by what we are learning from the sciences of human agency.

We can group such threats into two categories. One kind of threat purports to ‘undercut’ or ‘undermine’ our reasons for belief in free will. To develop a successful objection of this kind one must first identify the basis on which we believe in free will, and then show that this basis is unlikely to yield true beliefs. Although they are not without interest, undercutting objections have not been at the heart of the contemporary case for free will scepticism. Instead, those who invoke the sciences of agency to motivate free will scepticism typically mount rebutting objections to free will. Whereas undercutting arguments attempt to undermine our evidence for free will, rebutting arguments provide what are alleged to be positive reasons against free will.

<sup>1</sup> I am very grateful to Neil Levy, Richard Swinburne and an anonymous reviewer for their helpful comments on earlier versions of this paper.