

Doing without deliberation: automatism, automaticity, and moral accountability

NEIL LEVY¹ & TIM BAYNE²

¹Centre for Applied Philosophy and Public Ethics, University of Melbourne, Parkville, 3010 Australia & ²Department of Philosophy, Macquarie University, NSW 2109 Australia

Summary

Actions performed in a state of automatism are not subject to moral evaluation, while automatic actions often are. Is the asymmetry between automatistic and automatic actions justified? In order to answer this question we need a model of moral accountability that does justice to our intuitions about a range of modes of agency, both pathological and non-pathological. Our aim in this paper is to lay the groundwork for the development of such a model.

In all of us, even in good men, there is a lawless, wild-beast nature, which peers out in sleep.

Plato, The Republic

Introduction

In the early morning of 24 May 1987, Ken Parks drove 23 kilometres from his home in Pickering, Ontario to his in-law's house. He entered the house using the key he had brought with him, strangled his father-in-law unconscious, and fatally stabbed his mother-in-law. Parks then 'came to' and drove to a nearby police station where he reported his actions to the police, saying that he 'thought he might have killed some people' (Broughton *et al.*, 1994). Parks was charged with the murder of his mother-in-law, but acquitted on the grounds that he had committed his horrific actions while in a state of somnambulism (sleep-walking), and thus qualified for the defence of automatism.

The legal defence of automatism is well established (Fenwick, 1990; McSherry, 1998; Ridgway, 1996; Schopp, 1991), and reflects the commonsense judgment that individuals are not fully accountable for what they do in such states. But although it seems clear that automatism exculpates, it is not clear *why* it exculpates. What exactly is it about Parks's mental state that renders him non-culpable for his actions?

In an engaging study of moral responsibility in automatism, Robert Schopp (1991) argues that the automaton is not responsible for what she does because she is unaware of how the contemplated act

conflicts with her other desires, and such morally relevant information as her self-image and her moral beliefs. Although the automaton's wants and beliefs may cause her actions, they do not cause them 'in the manner characteristic of ordinary human activity' (Schopp, 1991, p. 145), and as a result her actions are not deeply attributable to her.

Schopp's account seems promising, but it generates the following problem. In some sense there is no such thing as 'ordinary human activity'. Instead, human activities involve a range of modes of agency, from the deliberative and reflective to the automatic and non-intentional. While the contrast between deliberative agency and automatisms seems clear, the contrast between automatisms and more automatic, non-deliberative, forms of 'ordinary human activity' is anything but clear (indeed a number of authors argue that automaticity is the norm for human activity [Bargh, 1997; Bargh & Chartrand, 1999]). And this is problematic, for although automatism exculpates, it is less obvious that *automaticity* exculpates. Consider the following vignettes:

- 1. A child runs out in front of Tim's new Volvo, and in order to avoid hitting her he drives into the embankment, damaging his car. He does not deliberate over what he should do; he is barely aware of the child's presence before he acts.
- 2. Janice is walking down a dark alley late at night when she hears a cry for help. She responds automatically, running in the direction from which the voice seemed to come. If she had reflected she would have chosen a less risky course of action. But she did not: she simply reacted to the sounds of distress.

Correspondence to: Tim Bayne, Department of Philosophy, Macquarie University, NSW 2109, Australia. Tel: + 61 2 9850 8837; Fax: + 61 2 9850 8892; E-mail: tbayne@scmp.mq.edu.au

3. The man walking a few paces ahead of Neil drops a fifty-dollar bill from his wallet without realizing it. Before Neil has time to think about what he should do, he scoops up the money and pockets it. It seems to him that he did not *decide* to keep the money. He just acted.

Although these actions are non-deliberative, we are inclined to regard them as properly subject to moral evaluation. We think it is appropriate to apply what Strawson (1962) called the reactive attitudes—attitudes such as praise and blame—to Tim, Janice and Neil; Tim and Janice are praiseworthy for their actions, while Neil ought to be censured for pocketing the money that did not belong to him. Individuals are perhaps *less* accountable for these kinds of actions than for their deliberative actions, but they seem to be accountable for them in a way that automatons are not. Can this common-sense judgment be justified?

In order to answer this question, we need a much finer-grained understanding of human agency than philosophers have typically provided. To a first approximation, we can think of agency as taking the following four forms:

- 1. Deliberative agency. We exercise deliberative agency when we deliberate or reflect on what we ought to do, attempting to evaluate our reasons for action in the light of our values, convictions, and beliefs about the world. This kind of agency has rightly impressed philosophers—largely, we suspect, because deliberative agency seems to be uniquely human—but in fact only a small proportion of our actions involve much in the way of deliberation or reflection.
- 2. Conscious agency. Although non-deliberative, much of human agency is conscious: we are typically aware of what we are doing and why we are doing it. Conscious agency roughly coincides with 'willed', 'controlled' or 'voluntary' agency (see e.g., Jahanshahi & Frith, 1998; Shallice, 1988; Spence, 2001b; Perner, 2003), although none of these terms is unproblematic.
- 3. Automatic agency. Automatic agency involves an absence—or at least a reduction—of the experience of doing. The paradigm of such behaviour is the over-learned action. One is usually (fully) conscious when performing an over-learned action, but one is not conscious of the over-learned action itself. The stock example here is that of driving a car (Bargh & Chartrand, 1999; Palmeri, 2002; Wheatley & Wegner, 2001). For the experienced driver, driving along a familiar road demands little in the way of attention, effort, or monitoring. One does it automatically—as one says—with one's attention on other things.
- 4. Automatistic agency. 'Automatism' is a broad and fairly vague term for a class of conditions in which one acts without being fully conscious of what one is doing. We think it useful to distinguish

between two kinds of automatisms: global automatisms and local automatisms. Global automatisms involve a global disruption of consciousness; they occur in the context of somnambulism, epileptic fugue, drug-related and trance states, and temporal and frontal lobe seizures. Individuals in these states perform fairly complex actions in a 'robotic' manner. Their environmental awareness is limited, and they tend to be amnesic for their actions (Fenwick, 1990). What we call 'local automatisms', by contrast, involve only a disruption of consciousness and control over a particular kind of action. A person with a local automatism is fully conscious, but they experience no sense of agency over a particular complex and apparently voluntary action (Wegner, 2002). Automatic writing—in which an agent writes but has no experience of directing her hand—is one form of local automatism, another is the anarchic hand syndrome, in which a person's hand (usually the right hand) engages in apparently purposive behaviour over which the agent has little or no direct control (Spence, 2001a; Spence, 2002).

Are these modes of agency discrete, clearly differentiated states, or are they four points on a continuum? Arguably Janice and Neil engage in conscious (but non-deliberative) agency, but perhaps Tim's actions are more automatic than conscious. Dual control accounts of agency (e.g., Norman & Shallice, 1986; Perner, 2003) seem to assume that the distinction between conscious and automatic processes is a clean one, but there is little reason to endorse this assumption. And even if conscious and automatic processes are distinct, complex actions involve nested hierarchies of both conscious and automatic processes: one can be conscious of certain high-level descriptions of an action but unconscious of the detailed motor routines involved in implementing it (Jeannerod, 1997). The important point for our purposes is not whether or not the divisions between these four modes of agency are hard and fast, but whether the differences between these modes of agency-such as they are-also mark differences in moral accountability. There is a prima facie case for thinking that they do: generally speaking, 'premeditated' homicide—murder 'in cold blood'—is worse than spur-of-the-moment killing, which is in turn is worse than killing performed as the result of an automatic action, which is in turn worse than killing while in a state of automatism. Are these judgments justified? And if so, what justifies them?

We will explore this question by examining three respects in which these modes of agency appear to differ: the degree to which they involve *deliberation*; the degree to which they reflect or manifest the agent's *character*; and the degree to which they involve *control* on the agent's part (Levy & Bayne, 2004).

Deliberation

How might deliberation and moral responsibility be related? One appealingly straightforward view is this: one is morally responsible only for instances of unimpaired deliberative agency. This view lines up with some of our intuitions concerning moral accountability. It accords with our view that children are less accountable than adults, for their deliberative abilities are less developed. It also accords with our attitudes to moral responsibility in the context of insanity and psychosis. A person suffering from psychosis might have deliberated over her actions at length, but since her deliberative abilities are impaired, she is not fully responsible for her actions.

This view also accords with our practice of not holding people accountable for their automatistic actions. There appears to be little room for deliberation in automatism. Whatever exactly the impairment of consciousness in global automatism involves-whether automatistic agents are completely unconscious; are conscious of less; or have a kind of occluded consciousness-global automatism offers an agent little opportunity for deliberation. The same point applies, for different reasons, to local automatisms. Consider the anarchic hand syndrome. Although individuals with this disorder engage in apparently purposive behaviour with their anarchic hand—they reach for a cup, open doors, and so on-these actions do not derive from a process of deliberation; they form part of no larger plan of action. Individuals with local automatisms are capable of deliberation, but their automatistic actions are not governed by the results of these deliberations, and hence are not deliberate.

But although the view that we are accountable only for deliberative actions may make sense of our attitudes to automatistic actions, it does not sit easily with a number of our other attitudes to accountability and agency. Bernard Williams (1982) and many other ethicists have suggested that there are situations in which spontaneous right action can be more praiseworthy than pre-meditated action: a person of good character would simply see what the appropriate thing to do is without needing to deliberate about it. (Note that there is a puzzling asymmetry between good action and bad action: while it is better to perform a good action spontaneously than on the basis of deliberation, deliberate wrongdoing always seems worse than spontaneous wrongdoing). Williams's view suggests that we can be responsible for non-deliberative actions.

This seems right. Recall again the three cases we introduced above. Although neither Tim, Janice, nor Neil deliberated before they acted, we want to hold them accountable for their actions. The mere fact that they failed to deliberate appears not to exculpate them. Consider the tragic case of the young father who simply forgot to drop his infant daughter at childcare, and instead left her to die in the parking lot

while he worked (Dennett, 2003). Though we blame him much less for his lapse than if he had deliberately set out to harm her, the mere fact that he neglected to deliberate adequately appears not to excuse him of all responsibility for his actions.

We might attempt to address this problem by suggesting that it is not deliberation per se that is relevant here but the opportunity to deliberate. Perhaps we can be accountable for the exercise of non-deliberative agency if we have had the opportunity to deliberate and failed to exercise it. Arguably, this is what happens in Neil's case: Neil had the opportunity to deliberate before pocketing the \$50 note but he failed to take it up. But although it is certainly true that we can be held responsible for not deliberating, we cannot account for all our intuitions regarding non-deliberative agency in this way. Neil is not only accountable for not deliberating, he is also accountable for doing what he did. (Following Williams, we might regard it as a failing of Neil's that he should need to deliberate here: a good moral agent would do the right thing without thinking about it). A second, and perhaps more serious, objection is that there are instances of morally accountable non-deliberative agency in which there is no opportunity for deliberation. Consider Tim, who has to decide what to do when a child runs out in front of his car. Even if Tim could have deliberated—which is perhaps doubtful—the costs of doing so would have been prohibitive; a rational agent ought not to deliberate in such situations, for by the time he has deliberated the opportunity for effective action will have passed.

One might attempt to handle this objection by appealing to the notion of backtracking. Perhaps our responsibility for automatic actions tracks back to earlier decisions—which may or at least ought to have involved deliberation—to perform actions, which led to these actions becoming habitual. In allowing ourselves to inculcate a habitual action, we ought to take into account the full range of circumstances in which that reflex might be triggered. If we inculcate a habit recklessly, and the reflex is triggered in a manner that causes harm (say), we are culpable despite the fact that we were not able to inhibit the response or to deliberate before we acted. Our responsibility tracks back to earlier moments when we could deliberate and control our actions.

While it is certainly true that we can be held accountable for recklessly developing habits and automatic patterns of behaviour, a backtracking account of this kind fails to deliver the general account for which we are looking. For one thing, it is difficult to see how it could deal with automatic behaviours that are hardwired rather than acquired. Tim's action, for example, does not seem to track back to earlier instances in which he engaged—or failed to engage—in deliberative agency.

A final point to consider: even if deliberation—or the opportunity for it—were to correlate perfectly

with moral accountability, we would want to know why this was so. One possibility is that deliberation is important in its own right. This position seems implausible to us. More plausible is the view that deliberative agency (or the possibility thereof) grounds accountability because it reflects or manifests the agent's character more fully than non-deliberative agency. A second possibility is that deliberation grounds accountability because it involves a greater degree of control over one's actions than non-deliberative agency. We turn now to explore these ideas.

Character

Perhaps we are responsible for an action to the extent that it reflects our character (Reznek, 1997). Since deliberate, pre-meditated actions reflect our character to a greater extent than non-deliberate actions, we are more accountable for them. Deciding to defraud one's employer after a protracted period of weighing up the pros and cons reflects more badly on one's character than seeing an opportunity for fraud and spontaneously pursuing it.

This approach to moral accountability also seems to make *prima facie* sense of our attitudes to automatism. Automatistic actions are often out of character. The Ken Parks case is an excellent example of this: Parks had no history of physical violence, and by all accounts he got on very well with his parents-in-law. His actions seem to be entirely out of character for him (Broughton *et al.*, 1994). Indeed, this seems to be true of many cases of violence in states of automatism (see Mahowald *et al.*, 1990). Perhaps this explains why those individuals who commit such acts are not responsible for them.

But despite its *prima facie* appeal, the attempt to root moral accountability in character faces problems. One problem concerns the assumption that non-deliberative agency does not reflect one's character. Consider Janice, who placed herself in some danger by running to the aid of a stranger. It is plausible to suppose that Janice would not have gone to the stranger's aid if she had had time to deliberate about her action. She might even say that she 'wasn't herself' when she acted-she doesn't regard herself as the sort of person who foolishly endangers herself in this way. But perhaps Janice is mistaken in making such claims. Perhaps she is wrong about the sort of person she is, and her non-deliberative actions reflect her character more accurately than do her considered judgments about what she thinks she would have done if she were in her right mind. (Of course, her considered judgments about what she thinks she would have done say something about Janice's character—they demonstrate that she is self-deceived in certain ways). Allowing that our non-deliberative actions can reflect our characters has the welcome implication that we can discover who we are from observing our actions (Stephens & Graham, 1996).

Although this line of thought has some appeal, it is not clear how far we ought to take it. Consider action in the context of dreaming—a cognitive state which seems to resemble global automatisms in a number of ways. St. Augustine wondered whether it is possible to sin while dreaming. As he rightly discerned, the answer is 'no', but is this because our dream actions and behaviour fail to reflect or manifest our characters (see Flanagan, 2000)? Janice might find herself going to the aid of a stranger in defiance of her own standing policy, but she would never find herself killing someone (unless in self-defence, or a result of a psychopathology) actions in which she may well engage while dreaming. When Janice becomes aware of her acts of dream violence, does she learn that she really does have violent tendencies lurking in the depths of her unconscious, or are these actions and motives not really attributable to her?

It might be said that we should not equate a person's character—their identity—with all of their beliefs, desires and various intentional states (some of which they might not know about). Perhaps we should not be identified with our non-endorsed desires and attitudes—states that we wish we didn't have and have attempted to expunge (Frankfurt, 1987). There is a sense in which non-endorsed states are less our own than endorsed states-states that reflect our all-things-considered judgments and with which we ourselves identify (Velleman, 1992)—but this point needs to be handled with some care, for it is also true that denying ownership of non-endorsed states can easily slide into a form of self-deception. Perhaps we should regard the products of automatic (and automatistic) agency as manifesting our characters no less than the products of conscious and deliberative agency. We return to this issue below.

Let us move now from the question of when an action might reflect one's character to the question of how character and moral accountability are related. Consider a case of global automatism in which an agent's actions *are* consonant with his character. Suppose, contrary to fact, that Parks did have a history of violence and aggression (see McSherry, 1998 for such a case). Would we hold such an individual accountable for their actions? We think not: the mere fact of automatism seems to remove—or at least greatly diminish—an agent's moral accountability.

Rather than the out-of-character nature of his actions *excusing* him, it may be the out-of-character nature of his actions that gives us reason to think that Parks lacked an important form of control over his behaviour, and it is this lack of control that removes Parks from the ranks of the morally accountable. The incongruity of his actions seems to be evidence that his behaviour was not governed or monitored by

executive decision-making processes. On this view, what's really doing the moral work here is not the link between character and agency but the link between control and agency; character is merely a heuristic for getting at control. Control, and not character, might therefore be the crucial element excusing automatistic actions.

Control

A common theme in discussion of moral responsibility is that responsibility and control are closely related (see Fischer & Ravizza, 1998). Roughly, one is responsible for an action to the degree that it is under one's control. Perhaps we can account for the differences between these modes of agency in terms of the degree to which they involve impairments in the agent's control over their behaviour.

It is useful to begin here with some distinctions. One distinction is that between *personal* and *sub-personal* levels of control. Personal control is the kind of control possessed by persons as such: here we think of the agent herself as directing and guiding her behaviour. Sub-personal control, by contrast, is possessed by sub-personal mechanisms. It is sub-personal control, for example, that allows one to track unexpectedly moving targets by hand prior to conscious awareness of the movement (Castiello, Paulignan & Jeannerod, 1991).

It is clearly personal control that is relevant to the discussion of moral accountability, for it is the person—rather than their sub-personal processes—with whom accountability rests. So if we could, say, align conscious and deliberative agency with personal control, and automatistic (and perhaps automatic) agency with sub-personal control, then we would have a tidy control-based vindication of our common-sense intuitions. Unfortunately, life is rarely tidy. Although few would challenge the claim that conscious and deliberative agency involve personal control, it is not clear that automatistic and automatic control can be parcelled out to sub-personal processes.

A second distinction that is useful here is the distinction between control itself and the *phenomenology* of control. Deliberative and conscious modes of agency involve the feeling of doing (Wegner, 2002). This is the sense that one is the doer of one's deeds and author of one's actions. The loss of the phenomenology of agency is a central and defining feature of local automatisms, both induced (automatic writing) and organic (the anarchic hand syndrome). The phenomenology of agency might also be lost in global automatisms, although it is difficult to tell.

What is the relationship between control and the phenomenology of control? Can the two come apart, and if so, which is it that grounds moral accountability: the feeling of control or control itself? Let's

start with the question of whether the loss of the phenomenology of control involves a loss of control itself. This is a difficult question to answer in full, but it seems clear that certain levels of control can survive in the absence of the phenomenology of control. Automatic writing is a case in point here. Individuals engaged in this practice have lost the sense of agency for what they are doing with their writing hand, but what they are doing with the hand is in some sense generated by their intentions, as can be seen by comparing automatic writing with the anarchic hand syndrome. Unlike automatic writers, individuals with an anarchic hand describe the hand as having a mind of its own, and are often forced to use their good hand to hold their bad hand down (Goldberg, Mayer & Toglia, 1981; Feinberg et al., 1992; Banks et al., 1989).

But despite some independence between control and the phenomenology thereof, it is plausible to suppose that certain kinds of control depend on the phenomenology of agency. Arguably, deliberation demands the feeling of doing (or at least the feeling that one can do). However, the precise nature of the link between control itself and the phenomenology of control remains mysterious, and without denying the importance of the phenomenology of control for accounts of moral responsibility we will leave it aside for the present and focus our attention on control itself. (It seems clear that the feeling of control cannot by itself ground accountability, for one can erroneously feel as though one is doing things that are being done by someone else. Nevertheless, it may be that the feeling of control is necessary for moral accountability).

There seems to be a tight connection between control and conscious and deliberative agency. It is not for nothing that conscious processes are often called controlled processes. While this label might be misleading in that it suggests that automatic processes are uncontrolled, it is certainly true that conscious agency involves a level of control not (typically) possessed by automatic and automatistic agency. Consciousness enables us to inhibit or veto our initial impulses to act. Absent consciousness, behaviour is guided by (fairly) modular actionroutines and over-learned scripts: control is local rather than global, and (generally) sub-personal rather than personal.

So the notion of control provides us with a morally relevant distinction between conscious and deliberative agency, on the one hand, and automatic and automatistic agency on the other. Does it also provide us with a morally relevant distinction between automatic and automatistic agency? We think it might.

The similarities between automatic agency and automatistic agency appear to run fairly deep, but the differences between them are important and must not be overlooked. Automatic agency is sensitive to disruption and resistance in the way that automatistic

agency is not. The disruption of automatic agency generally leads to the resumption of conscious control over the action. When we meet with unexpected road conditions we switch from automatic to conscious agency, and are immediately conscious of what we are doing. Although automatic actions are (by definition) not directly monitored, they are indirectly monitored to the extent that the agent remains ready to move to conscious forms of agency should that be necessary.

In contrast, automatistic agency appears to be resistant to disruption. There is reason to think that aggression in automatism might be due to the frustration of an action, which appears to be 'essentially pre-programmed' (Broughton *et al.*, 1994). Whereas the frustration of automatic agency typically leads to the re-establishment of executive control, the frustration of action in instances of global automatism leaves the individual confused and without the resources of conscious control. Automatic agency involves a form of monitoring that appears to be absent in global automatism.

As an aside, we can note that this line of thought supports the parallel between global automatism actions and normal dreaming. (Lucid dreaming might be importantly different from normal dreaming here.) Automatism shares many features with dreaming. There is reduced awareness of the environment (greater awareness in automatism than in ordinary dreaming, but perhaps awareness of the environment is not entirely lacking in ordinary dreaming consciousness), and frequently partial or total amnesia of the dream-state upon awakening. Somnambulism (with some rare exceptions due to brain lesions) occurs during non-REM sleep, which was formerly believed to be dreamless, but which we now know to have its own characteristic dream mentation (Flanagan, 2000). Perhaps global automatistic actions are the product of a state of consciousness akin to that of non-REM dreaming, in which action plans are initiated and carried out with little awareness of the environment and diminished understanding of the nature and significance of the

Let us pull the various threads of this discussion together. We suggest that automatism excuses, *interalia*, because the agent: initiates actions which she would not normally perform; because she lacks executive control over her actions; and because frustration of her action-plans fails to re-establish conscious control, as in automatic non-deliberative action, but leaves sub-personal mechanisms in control.

This account of responsibility for automatistic actions combines elements of both the control and character models of accountability. Agents who act automatistically are not responsible for their actions because they cannot exert the right kind of control over their actions, where the right kind of control is control that manifests character in a deep sense.

Those sub-personal mechanisms that initiate and guide behaviour do not reflect the agent's endorsed values in such a way as to make it the case that they are accountable for them.

To say this is not to provide a model of moral accountability that does justice to the full range of complex (and possibly confused) intuitions we have about responsibility in non-deliberative agency. Such a project is too ambitious for the present context. Instead, we have pointed to some of the basic components available for the construction of such a model, and we have drawn attention to some of the difficulties that confront one in attempting to put these components together in a satisfactory way.

Acknowledgements

We are very grateful to Matthew Broome, George Graham, Melissa Green, Elisabeth Pacherie, Sean Spence and Dan Wegner for discussions related to this paper.

References

Banks, G., Short, P., Martinez, J., et al. (1989). The alien hand syndrome: clinical and post mortem findings. Archives of Neurology, 46, 456–459.

BARGH, J.A. (1997). The automaticity of everyday life, In R.S. Wyer, Jr. (Ed.), *The automaticity of everyday life, advances in social cognition*, X (pp. 1–62). Mahwah, NJ: Erlbaum.

BARGH, J.A. & CHARTRAND, T.L. (1999). The unbearable automaticity of being. *American Psychologist*, 54, 462–479.

Broughton, R., Billings, R., Cartwright, R., Doucette, D., Edmeads, J., Edwards, M., et al. (1994). Homicidal somnambulism: a case report. Sleep, 17, 253–264.

CASTIELLO, U., PAULIGNAN, Y. & JEANNEROD, M. (1991). Temporal dissociation of motor responses and subjective awareness: a study in normal subjects. *Brain*, 114, 2639–2655.

Dennett, D. (2003). Freedom evolves. New York: Viking. Doris, J. (2001). Lack of character: personality and moral behavior. Cambridge: Cambridge University Press.

Feinberg, T.E., Schindler, R.J., Flanagan, N.G., et al. (1992). Two alien hand syndromes. *Neurology*, 42, 19–24.

Fenwick, P. (1990). Automatism, medicine and the law. *Psychological Medicine Monograph*, 17, 1–27.

FISCHER, J.M. & RAVIZZA, M. (1998). Responsibility and control. Cambridge: Cambridge University Press.

FLANAGAN, O. (2000). *Dreaming souls*. Oxford: Oxford University Press.

Frankfurt, H. (1987). Identification and externality. In H. Frankfurt (Ed.), *The importance of what we care about* (pp. 58–68). Cambridge: Cambridge University Press.

GOLDBERG, G., MAYER, N.H. & TOGLIA, J.U. (1981). Medial frontal cortex and the alien hand sign. Archives of Neurology, 38, 683–686.

JAHANSHAHI, M. & FRITH, C.D. (1998). Willed action and its impairments. Cognitive Neuropsychology, 15, 483–533.

Jeannnerod, M. (1997). The cognitive neuroscience of action. Oxford: Blackwell.

- Levy, N. & Bayne, T. (2004). A will of one's own: Consciousness, control and character. *International Journal of Law and Psychiatry* (in press).
- Mahowald, M.W., Bundlie, S.R., Hurwitz, T.D. & Schenck, C.H. (1990). Sleep violence: forensic science implications: polygraphic and video documentation. *Journal of Forensic Science*, 35, 413–432.
- McSherry, B. (1998). Getting away with murder: dissociative states and criminal responsibility, *International Journal of Law and Psychiatry*, 21, 163–176.
- NORMAN, D.A. & SHALLICE, T. (1986). Attention to action: willed and automatic control of behavior. In R.J. DAVIDSON, G.E. SCHWARTZ & D. SHAPIRO (Eds.), Consciousness and self-regulation (Volume 4, pp. 1–18). New York: Plenum.
- Palmeri, T.J. (2002). Automaticity, In L. Nadel (Ed.), Encyclopedia of cognitive science (pp. 290–301). London: Nature Publishing Group.
- Perner, J. (2003). Dual control and the causal theory of action: the case of non-intentional action. In N. EILAN & J. ROESSLER (Eds.), *Agency and self-awareness* (pp. 218–243). Oxford: Oxford University Press.
- REZNEK, L. (1997). Evil or ill? Justifying the insanity defence. London: Routledge.
- RIDGWAY, P. (1996). Sleepwalking—insanity or automatism? E Law: Murdoch University Electronic Journal of Law, 3.
- Schopp, R.F. (1991). Automatism, insanity, and the psychology of criminal responsibility: a philosophical inquiry. Cambridge: Cambridge University Press.

- SHALLICE, T. (1988). From neuropsychology to mental structure. Cambridge: Cambridge University Press.
- Spence, S. (2001a). Alien control: from phenomenology to cognitive neurobiology. *Philosophy, Psychiatry, Psychology*, 8, 163–172.
- Spence, S. (2001b). Disorders of willed action. In P. Halligan, C. Bass, & J. Marshall (Eds.), Contemporary approaches to the study of hysteria. Oxford: Oxford University Press.
- Spence, S. (2002). Alien motor phenomena: a window on to agency. In S. Spence & P.W. Halligan (Eds.), *Pathologies of body, self and space* (pp. 211–220). Hove, East Sussex: Psychology Press.
- STEPHENS, G.L. & GRAHAM, G. (1996). Psychopathology, freedom, and the experience of externality. *Philosophical Topics*, *24*, 159–182.
- STRAWSON, P.F. (1962). Freedom and resentment. Proceedings of the British Academy, 48, 187–211.
- Velleman, D.J. (1992). What happens when someone acts? *Mind*, 101, 462–481.
- WEGNER, D. (2002). The illusion of conscious will. Cambridge, MA: MIT Press.
- WHEATLEY, T.P. & WEGNER, D.M. (2001). Automaticity in action. In N. J. SMELSER & P. B BALTES (Eds.), *International encyclopedia of the social and behavioral sciences* (pp. 991–993). London: Pergamon.
- WILLIAMS, B. (1982). Persons, character and morality. In B. WILLIAMS (Ed.), *Moral luck* (pp. 1–19). Cambridge: Cambridge University Press.