



# The Phenomenology of Agency

Tim Bayne\*

University of Oxford

---

## Abstract

This paper provides an overview of recent discussions of the phenomenology of agency. By ‘the phenomenology of agency’ I mean those phenomenal states that are associated with first-person agency. I call such states ‘agentive experiences’. After briefly defending the claim that there is a phenomenology distinctive of first-person agency, I focus on two questions: (i) What is the structure of agentive experience? (ii) What is the representational content of agentive experience? I conclude with a brief examination of how agentive experiences might be generated and what role they might play in the subject’s cognitive economy.

---

The motion of our body follows upon the command of our will. Of this we are at every moment conscious.

David Hume, *Enquiry Concerning Human Understanding*

## 1. Introduction

Until recently the phenomenology of agency was rather neglected, overlooked by both philosophers of action and philosophers of consciousness alike. Thankfully, all that has changed, and of late there has been an explosion of interest in what it is like to be an agent.<sup>1</sup> This burgeoning discussion straddles the traditional boundaries between disciplines: philosophers of psychopathology are speculating about the role that unusual experiences of agency might play in accounting for disorders of thought and action; cognitive scientists are developing models of how the phenomenology of agency is generated; and philosophers of mind are drawing connections between the phenomenology of agency and the nature of introspection, phenomenal character, and agency itself. My aim in this paper is not to provide an exhaustive survey of this recent literature, but to provide a guide for the perplexed — a map by means of which newcomers to the debate might orient themselves. The reader should be warned that although my primary goal is to describe positions rather than defend or dismantle them, I do have a horse in some of these races.

## 2. Identifying the Target Domain

What is the phenomenology of agency? There is no simple answer to this question. In the widest sense one might take ‘the phenomenology of agency’ to refer to any phenomenal state that is associated with agency. I will employ on a narrower sense of the term, focusing only on the phenomenology of first-person agency. Even so, the term encompass a wide range of phenomenal states. Consider the following vignette:

It’s 4 a.m. You awaken to the smell of smoke inside your apartment. You stumble out of bed and reach for the door so as to leave the room and escape the fire. Something is leaning against the door, and you must strain in order to open it. Finally, the door begins to move — you experience yourself as opening the door, and moving into the corridor. Which way should you turn? Right or left? Suddenly, you remember that the only copy of your half-finished manuscript is in the study and at the mercy of the flames. You deliberate for a second: should you go back for the manuscript or continue down the corridor and out of harm’s way? You decide to risk it — life without the manuscript would not be worth living.

This vignette contains reference to multiple examples of agentic self-awareness. You experience the movement of your arm towards the door as involving your own agency — as something that you are doing rather than as something that is happening to you. In reaching for the door you experience yourself as doing something in particular, namely, reaching for the door. You experience yourself as exerting effort in pushing against the door. You experience your decision to head right (rather than left) as autonomous, as ‘up to you’. You experience yourself as deliberating — as ‘making up your mind’ — about whether or not to go back for the manuscript.

Whether or not one takes ‘the phenomenology of agency’ to include all the above-mentioned states depends, in large part, on one’s view of cognitive phenomenology. Some theorists allow that cognitive states can enjoy proprietary phenomenal character (see e.g. Pitt; Siewert; Strawson, *Mental Reality*); others do not (see e.g. Carruthers, ‘Conscious Experience’; Lormand; Tye). Those who recognise cognitive phenomenology are likely to allow that agentic phenomenology might include representations of what one is doing (or at least trying to do), although such theorists might not be entirely happy about calling such states ‘experiences’. Those who reject the possibility of cognitive phenomenology will restrict the scope of agentic phenomenology to ‘low-level’ states with non-conceptual content. The issues raised by cognitive phenomenology are too complex to tackle here, and for the most part I will restrict my attention to those forms of the phenomenology of agency that might be recognised as such even by those who deny that cognitive states possess phenomenal character.<sup>2</sup> I would include on such a list the experience of a movement (or a mental event) as: an action (of one’s own); an action that one is in control of; an

action that one is performing with a certain degree of effort; and as an action that one is performing freely. Note that I use ‘experience’ non-factively (that is, intensionally), such that one can experience that things are thus-and-so even when they are not thus-and-so.

### 3. *Scepticism*

Some are sceptical of the claim that there is a phenomenology of agency. Thoroughgoing sceptics deny that there is anything at all that it is like to be an agent; they deny that ‘the experience of acting’, ‘the experience of deliberating’, ‘the experience of trying’ and so on are referentially successful. Scepticism of this kind is rarely defended in print, but it is frequently encountered in conversation.

We can distinguish two argumentative strategies for scepticism. The first, the Humean strategy, begins with introspection. The Humean sceptic ‘looks inside’ and fails to find anything to which ‘the experience of acting’, or ‘the experience of being an agent’ might correspond. She then argues that since there is no reason to think that she might be uniquely lacking in agentic experience, it is doubtful that anyone enjoys experiential states to which these phrases might refer.

One challenge to the Humean strategy calls into question the assumption of phenomenal uniformity on which the argument rests. Although we tend to assume that human beings enjoy pretty much the same range of phenomenal states, it is entirely possible that this assumption is mistaken (or at least overblown) — after all, we know that there can be striking individual differences with respect to certain kinds of phenomenal states, such as imagery experiences. Alternatively, one might argue that the sceptic is overlooking features of her own experience. The advocate of agentic experience need not hold that agentic experiences are phenomenologically vivid or easy to discern; indeed, it is common for agentic experience to be described as recessive — as typically confined to the margins of consciousness.

A second argument for scepticism involves the claim that agentic content cannot be experientially encoded — it is not ‘experientially admissible’. According to Christine Korsgaard, ‘to experience something is (in part) to be passively receptive to it, and therefore we cannot have experiences of activity as such’ (204). I find this claim puzzling: why should the mere passivity of experience prevent agentic content from being experientially encoded? After all, belief is (typically) thought to be passive, but there is no doubt that we can have beliefs about activity as such. More plausible is the claim that the architecture of consciousness might prevent agentic content from being experientially encoded. For example, one might identify phenomenal consciousness with the output of Fodorian (1983) perceptual modules. This approach to the architecture of phenomenal consciousness would seem to be at odds with the claim that there is a

distinctive phenomenology of first-person agency, for it is far from obvious that agentive experiences could be generated by Fodorian modules.

Although more plausible than its predecessor, it seems to me that this argument for scepticism can also be disarmed. For one thing, there is some reason to suspect that at least some components of agentive experience are subserved by low-level modular systems (see section 11). These modules might be output (action) modules rather than input (perception) modules, but they are modules nonetheless. More generally, the objection is vulnerable insofar as there are independent reasons to resist the assumption that phenomenal consciousness is restricted to modules. Not only is there much to be said in favour of a proprietary phenomenology of cognition, there is also 'fringe phenomenology' to reckon with: 'tip of the tongue' experiences, *déjà* and *jamais vu* experiences, and feelings of recognition, familiarity and significance are just some of the many experiential states that are not obviously produced by Fodorian modules.

The central challenge facing the sceptic is to explain (away) the apparent existence of agentive experience. The experience of oneself as an agent is, for many of us, a robust and readily identifiable phenomenon. The sceptic needs to explain why it is so plausible to *think* that there is a phenomenology of agency even though there isn't. Perhaps, the sceptic might suggest, so-called agentive experiences are actually more familiar experiential states that are misdescribed as distinctively agentive. What might these familiar states be? I can think of two candidates: they could be bodily sensations of various kinds, such as experiences of movement (or the potential for movement), or they could be conscious judgements whose contents concern agency. Let's examine these two proposals in turn.<sup>3</sup>

Although agentive experiences are intimately associated with bodily sensations of various kinds, we should resist identifying agentive experiences with bodily sensations. It is one thing to experience one's body as moving (being about to move; being capable of moving), but it is quite another thing to experience oneself as acting (being about to act; being able to act). Such commonplace phenomena as involuntary muscular spasms reveal the gulf between experiences of movement and experiences of agency. The attempt to reduce agentive experience to the experience of bodily movement also struggles to account for the experience of mental agency. Although deliberation, decision-making, mental effort and the voluntary allocation of attention presumably involve movements in the head, this is not how these states are experienced.

What about the attempt to reduce agentive experience to conscious judgements about agency? Obviously, one cannot explain away the appearance of agentive experience by appeal to cognitive phenomenology unless there is such a thing as cognitive phenomenology, and as we noted above, the existence of cognitive phenomenology is not uncontroversial. But even those who are attracted to cognitive phenomenology should resist the suggestion that agentive experience is purely doxastic in nature. Just as the Müller-Lyer

illusion shows that the contents of visual experience can be at odds with those of judgement, so too certain disorders of agency suggest that the contents of agentic experience can also be at odds with those of judgement.

Consider two disorders of agency: the anarchic hand syndrome (Della Sala et al.; Goldberg and Bloom) and utilization syndrome (Lhermitte; Estlinger et al.). The two syndromes are similar in that each involves an inability to inhibit stimulus-driven actions. The patient with an anarchic hand (or should we say 'the patient's anarchic hand?') will take food from another's plate and undo buttons that have only just been done up; the patient with utilization behaviour will put on multiple pairs of sunglasses, even when she is already wearing sunglasses. These actions need not conform to the patient's over-arching goals, and — at least in the case of the anarchic hand syndrome — may even be at *odds* with the patient's goals. (The patient doesn't want to take food from his neighbour's plate.) But despite their behavioural commonalities, these two disorders give rise to very different reports: whereas patients with utilization behaviour show no inclination to disown their actions, patients with an anarchic hand typically describe the hand as 'having a will of its own'. This difference is vividly illustrated by a case in which a patient exhibited utilization behaviour with his right hand and anarchic hand behaviour with his left hand: he was unconcerned about the former but troubled by the latter (Marcel)!

Although phenomenologically detailed descriptions of these disorders are hard to come by, it is not implausible to suppose that the agentic experience of anarchic hand patients differs from that of patients with utilization behaviour. It is because the patient with an anarchic hand fails to experience himself as the agent of the movements of his anarchic hand that he denies authoring its movements, and it is because the patient with utilization behaviour has normal experiences of agency (or at least, does not have abnormal experiences of agency) that she fails to disown her actions. But it is not hard to imagine that one might convince the patient with an anarchic hand that he is acting even though he doesn't experience himself as acting, nor is it hard to imagine that one might convince the patient with utilization behaviour that she is not acting despite the fact that she experiences herself as acting (or doesn't experience herself as not acting). A patient simultaneously manifesting an anarchic hand and utilization behaviour might believe that both (or neither) of the movements in question constitute actions of theirs, but none of this will change his or her agentic experience.

Of course, all of this is consistent with the claim that certain *components* of agentic self-awareness might involve (or simply be) judgments. Consider one's experience of what one is currently doing. It seems implausible to suppose that one could believe that one is performing one type of action (say, waving to a friend) by means of a particular physical movement whilst at the same time experiencing oneself as performing another type of action

(say, practicing a dance movement) by means of that same physical movement. One's awareness of what one is doing seems not to be subject to cognitive encapsulation, and hence might well qualify as a judgement of a certain kind.

#### 4. *The Structure of Agentive Experience*

Let us assume that there is a proprietary phenomenology of agency. How should we conceptualise this phenomenology? There are two general issues to be addressed: one concerns the *structure* of agentive experience, the other concerns the *contents* of agentive experience. I examine the structure of agentive experience in sections 4 to 6 and the contents of agentive experience in sections 7 to 9.

What is the structure of agentive experience? One possibility is that agentive experiences are *descriptive* (or *thetic*). On this view, agentive experiences are in the business of saying how things are; they have veridicality conditions and a mind-to-world direction of fit. Descriptivism is of a piece with the trend towards representationalism in the analysis of phenomenal character generally. Roughly speaking, representationalists hold that the phenomenal character of a state (that is, what it is like to be in the state) can be analysed in terms of its representational content (that is, the state of affairs that it represents).<sup>4</sup>

Some theorists seem to assume that descriptivism is the only game in town. For example, Horgan et al. move directly from the claim that there is a phenomenology of agency to the question of what its veridicality conditions might be. But things are not quite that straightforward, for it is possible to hold that agentive experiences are intentional but not descriptive — that is, one could hold that agentive experiences have satisfaction conditions but not veridicality conditions. Searle (*Intentionality*) appears to hold such a view. According to Searle, experiences of acting say how the world should be rather than how it is. I will call this the *directive* (or *telic*) account of agentive experience. I return to the directive account shortly.

A third position in logical space holds that agentive experiences have both a mind-to-world *and* a world-to-mind direction of fit. To use Millikan's term, agentive experiences might be *pushmi-pullyu* representations (Millikan). The pushmi-pullyu account promises to do justice to the thought that agentive experiences both represent oneself as acting (as the descriptive account has it) and are also features of intentions themselves (as the directive account has it). Of course, the notion of a pushmi-pullyu representation is not unproblematic, for it is far from clear that a single state can have both a world-to-mind and a mind-to-world direction of fit.

Although the directive and pushmi-pullyu accounts flout the letter of representationalist analyses of phenomenal character, both approaches fall within its spirit, for they accept that the analysis of agentive phenomenology can be carried out with reference to intentional content. A fourth account of agentive experience takes issue with this assumption. According

to what we might call the ‘raw feels’ account, agentic experience involves raw phenomenal feels, non-intentional phenomenal properties. I know of no defence of this view, but those who are attracted to a raw feels account of sensations are likely to be attracted to a similar treatment of agentic experience. To my mind, the most convincing case for the raw feels account concerns the experience of effort. Consider what it’s like to complete a set of press-ups, attempt to solve an unsolvable puzzle, or resist temptation. One’s efforts to succeed in these endeavours can be assessed for success, but it is less clear that one’s *experiences* of effort can be similarly assessed.

So we have four accounts of agentic experience: the descriptive account, the directive account, the pushmi-pullyu account, and the ‘raw feels’ account. The last two accounts have not been developed in any detail, and I will leave them for others to pursue. In the next section, I examine Searle’s version of the directive approach.

### 5. *The Directive Account*

In his seminal discussion of agentic experience in *Intentionality*, Searle argued that at least one kind of agentic experience — ‘the experience of acting’ — is directive.<sup>5</sup> Searle contrasts experiences of acting with visual experiences.

As far as *Intentionality* is concerned, the differences between the visual experience and the experience of acting are in the direction of fit and in the direction of causation: the visual experience stands to the table in the mind-to-world direction of fit. If the table isn’t there, we say that I was mistaken, or was having a hallucination, or some such. And the direction of causation is from the object to the visual experience. If the Intentional component is satisfied it must be caused by the presence and features of the object. But in the case of the experience of acting, the Intention component has the world-to-mind direction of fit. If I have this experience but the event doesn’t occur we say such things as that I *failed* to raise my arm, and that I *tried* to raise my arm but did not succeed. And the direction of causation is from the experience of acting to the event. Where the Intentional content is satisfied, that is, where I actually succeed in raising my arm, the experience of acting causes the arm to go up. If it didn’t cause the arm to go up, but something else did, I didn’t raise my arm: it just went up for some other reason. (*Intentionality* 88, emphasis in original; see also 123f)

According to Searle, experiences of acting just are intentions of a certain kind, what he calls ‘intentions in action’. Searle distinguishes intentions in action from other intentions — prior intentions — in two ways: (i) some actions, such as spontaneously pacing around a room, occur without the formation of prior intentions, whereas all actions involve intentions in action; (ii) prior intentions cause actions, but intentions in action are components of actions (84).

What does Searle mean by ‘the experience of acting’? Does he mean to identify the experience of acting with the experience (as) of *having* a certain intention-in-action or with the experience (as) of *satisfying* an intention-in-action? Given that Searle grants that intentions in action are, roughly, tryings (‘Reply to Wakefield and Dreyfus’ 298), it seems to follow that the experience of acting must correspond to the experience of having an intention in action. But is this all there is to the experience of acting? Searle’s discussion of the experience of acting suggests not, for he makes much of William James’s example of the patient who is asked to raise his anesthetized arm. The patient’s eyes are closed, and unbeknown to him his arm is prevented from moving. Upon opening his eyes, the patient is surprised to discover that his arm has not moved. What is the content of the patient’s experience? Searle describes it as an experience of *trying* but *failing* to raise one’s hand (*Intentionality* 89). It seems to me that this description is inaccurate. Intuitively, the patient’s experience is one of *raising one’s hand* — he experiences himself as having just raised his arm. After all, the patient is surprised to discover that his hand is not raised. Why would he be surprised if he experienced himself as merely trying to raise his hand? Searle himself points to a parallel between experiences of acting and visual experiences without appearing to notice that the parallel is at odds with his own account (see 87). Searle’s account appears to entail that what it’s like to experience oneself as merely trying to do something is identical to what it’s like to experience oneself as actually doing something, and that seems highly implausible.<sup>6</sup>

But perhaps ‘the’ experience of acting actually has two components: an experience of trying to do something (which has directive structure), and an experience of actually doing that thing (which has descriptive structure). This proposal might be accepted by those, such as O’Shaughnessy, who identify experiences of trying with tryings themselves. Such an identification would rule out the possibility of a person experiencing themselves as trying to do something that they are not trying to do. (However, there would still be space for the possibility of someone trying to do something that they do not experience themselves as trying to do.) The descriptive account of trying, by contrast, must allow that it is possible — at least in principle — to experience oneself as trying to do something when one is not trying to do anything at all. Whether or not such an implication is a liability or an asset for the descriptive view is an open question.

## 6. *The Descriptive Account*

We can distinguish two central projects for the descriptivist. The first project is that of giving an account of the contents of agentive experience. Just how demanding this task is depends, to some degree, on the richness of agentive content. Some theorists subscribe to an *austere* conception of



agentive experience; we might call such theorists ‘phenomenal conservatives’. Phenomenal conservatives hold that the contents of agentive experience contrast sharply with those of other experiential modalities, such as vision. Whereas visual experience presents us with a richly tapestried world, replete with colours, shapes, sizes, movement, textures and (arguably) much more, the world that is presented to us in agentive experience is ‘bare’: agentive experience can represent oneself as acting, as being in control of one’s actions, and as acting with effort, and that’s roughly it. At the other end of the spectrum lie the phenomenal liberals, who hold a rich conception of agentive experience. Liberals do not necessarily take the contents of agentive experience to be as rich as those of vision, but they certainly think that the contents of agentive experience can — and typically do — exceed the mere representation of agency, control and effort. There are two ways to hold an expansive view of agentive experience. One might adopt liberalism because one thinks that agentive experiences is capable of encoding a rich array of agency-related properties and relations, or because one thinks that agentive experience can be directed towards a rich array of bodily (and perhaps mental) states and events. My sympathies lie towards the liberal end of the scale, but I will not attempt to argue for liberalism here.

The second project on the descriptivist agenda is that of determining how accurate agentive experiences are. We tend to suppose that agentive experience is generally reliable. We assume that when we experience ourselves as acting we are acting, and that when we experience ourselves as acting in a certain way — say, with effort, or freely — then we are indeed acting in the relevant way. But perhaps these assumptions are misplaced.

Let us distinguish two kinds of error to which agentive experience might be subject: *superficial* error and *deep* error. Superficial error occurs when the mechanisms that generate experiences of agency ‘misfire’. This is the sort of thing that occurs in the Jamesian vignette discussed earlier: the patient experiences himself as acting despite not moving. It would not be particularly surprising to discover that the mechanisms responsible for agentive experience are subject to error — after all, few representational systems are infallible. Deep error occurs when the mechanisms responsible for agentive experience produce non-veridical representations even when functioning normally and in normal environments. Some hold that the mechanisms of colour experience are systematically erroneous, in that even in normal conditions colour experiences represent objects as having properties that they lack. (Colours, some say, are experienced as simple, intrinsic properties of objects, but so-understood there are no colours.) Similarly, one might argue that agentive experience involves deep error to the extent that we experience ourselves as having (say) libertarian free will, even though no one has (or perhaps even could have) libertarian free will.

Of course, the question of how reliable agentive experience is cannot be divorced from the question of what its content includes. The more

that is packed into the content of agentive experience, the greater the threat of error ('the more you say, the more you can be wrong about'). In light of this, we might expect those who regard agentive experience as riddled with error as buying into a liberal conception of agentive content, and, correlatively, we might expect those who regard agentive experience as largely trustworthy to be sympathetic to a rather more conservative conception of agentive experience. Let us examine how these general issues play out with respect to three components of agentive experience: the experience of mental causation, the experience of the self, and the experience of free will.

### 7. *Mental Causation*

As folk, we appear to be deeply wedded to a conception of mental states according to which they not only rationalize our actions but also cause them, and their causal efficacy is dependent on their content. This conception of folk-psychology is not uncontested, but it is widely endorsed. Is it also a component of agentive experience?

A number of theorists — particularly those within the cognitive sciences — assume that it is. According to what we might call 'the causal account', the experience of purposive goal-directed behaviour involves (or perhaps just is) an experience of mental causation (Hohwy). In its simplest incarnation, the causal account says simply that to experience one's movements as one's actions involves experiencing them as caused by mental states (that one also experiences as one's own). This sketch can be embellished in various ways. For example, one might suggest that the experience of goal-directed agency involves the experience of causal relations between mental states. Consider again the vignette with which I opened this paper, in which you experience yourself as opening a door in order to leave your apartment and thus escape a fire. Arguably, this scenario includes an experience of causal relations as nested within each other: you experience your intention to flee the fire as causing your intention to leave the apartment, and in turn experience your intention to leave the apartment as causing your intention to open the door. If one accepts that the basic experience of agency includes an experience of one's intentions as causing one's movements, then there is at least a *prima facie* case for thinking that one can also experience some mental states as causing other mental states.

But what should we say about the basic causal proposal, namely, that the phenomenology of agency involves the experience of one's movements as caused by one's intentions? Some theorists reject the causal account on the grounds that causation cannot be — or at least *is not* — experienced. I don't have much sympathy with this criticism, for it seems evident to me that causal relations can be experientially encoded. (One sees the bat as *causing* the ball to ricochet away.) A more plausible objection holds that although causal relations can be experientially encoded, they are not

contained within the normal experience of autonomous agency. One might take the pathological quality of agentive experience in addiction and Tourette's syndrome to *consist in* the experience of mental causation. And — so the argument goes — if the experience of mental causation characterises pathological agentive experience, then it cannot also characterise non-pathological agentive experience. If successful, this line of argument would demonstrate that it is possible to experience one's own mental states as causally efficacious, but it would undermine the idea that the normal experience of autonomous agency can be understood in such terms.

Do pathological experiences of agency really involve the experience of mental causation? There is certainly reason to doubt this in the case of Tourette's syndrome, for those with the condition typically describe their tics as being under their control (Bliss; Cohen and Leckman). Addicts may describe themselves as feeling overwhelmed by their urges, but it is not entirely clear that this experience brings with it a sense of being caused to act by one's urges. Perhaps the phenomenology of addiction involves just the sense that one has no (or at least little) control over one's urges. Or perhaps we can characterise the experience of addictive compulsion in terms of the phenomenal strength of desire rather than its perceived causal efficacy.

Even if these cases do involve the experience of mental causation, it is far from clear that this is what accounts for their pathological nature. Perhaps the difference between the normal case and the pathological case is that in the normal case one experiences one's *intentions* as causally efficacious whereas in the pathological case one experiences one's *desires* (or emotions or urges) as causally efficacious. Or perhaps the difference resides in the fact that in the normal case the causally efficacious mental state is experienced as one's own, whereas in the pathological case it is experienced as alien in some way. Neither of these suggestions is obviously right, but both are worthy of serious consideration.

Arguably the most potent objection to the causal account is that it simply misrepresents the nature of agentive experience. O'Connor, Horgan et al., and Wakefield and Dreyfus confess that they are unable to discern anything that might answer to 'the experience of mental causation' in themselves. These authors grant that they experience their movements as *realising* their intentions and as *satisfying* their desires, but they deny that they experience their movements as caused by their mental states. (They claim, instead, to experience their movements as caused by, or at least having their source in, themselves.) I have some sympathy with this view. At the same time, it is not difficult to make the causal account attractive. Suppose that Malebranche is right, and that mental and physical events are not related by direct commerce but only by a sort of pre-established harmony: either by divine will or cosmic coincidence it just so happens that one's actions are in accord with one's agentive experiences, but there is no sense in which the objects of one's agentive experiences are causally implicated in

one's actions. It seems clear that there would be something misleading about the phenomenology of agency in such a world. Arguably, we wouldn't be so concerned to make the world safe for mental causation if the causal account of agentic experience weren't true.<sup>7</sup>

An interesting variant of the causal account — endorsed by both Searle (*Intentionality*) and Mossel — holds that experiences of agency represent *themselves* as causing the target movements. I find this suggestion puzzling. The experience I have of myself as raising my arm might cause me to believe that I am raising my arm, but it does not seem to me that *it* — that is, the agentic experience itself — causes my arm to move. Searle is drawn to the view that experiences of acting are causally self-referential because (as we have seen) he regards experiences of acting as intentions, and he regards intentions as representing their own causal efficacy. Obviously, this argument for the causally self-referential conception of agentic experience turns on a commitment to the directive account of agentic experience, and will have no grip on those who prefer alternative conceptions of agentic experience. In fact, one might be inclined to turn this argument on its head and take the fact that agentic experiences are not causally self-referential as a reason to reject the directive account of agentic experience.

### 8. *The Place of the Self*

A second focus of debate concerns the way in which the experience of oneself enters into the phenomenology of agency. Although most theorists agree that the self is represented in agentic experience in some way or other there is much disagreement about exactly how the self is agentially represented.

A deflationary account of the role of the self in agentic experience accords the self a 'place-holding' role in agentic experience. On this view, agentic experience brings with it a distinction between things that one does and things that happen to one, but that's all. A number of authors have argued that we need to go beyond deflationism in one (or both) of two ways. Some argue that agentic experience brings with it a particular conception of the relationship between the self and its acts. According to what we might call the 'agent causal account', the phenomenology of agency includes an experience of the self as a cause, where the causal role of the self is not to be understood as derivative on or reducible to the causal role of the self's mental states. Ginet, Horgan et al. and O'Connor each express some sympathy with this view, although they do not necessarily endorse agent causation as such.<sup>8</sup> (Note that the agent causal account, as I construe it here, does not demand the representation of oneself as an *unmoved mover*.)

Some have rejected the agent causal account on the grounds that the very notion of agent causation is incoherent (Searle, *Rationality in Action*

82). It is not at all clear that the notion of agent causation is incoherent, but even if it is it may not follow that the agent causal account is false, for it might be possible for experience to represent impossible states of affairs as such. (Consider what it is like to look at an Escher drawing.) Others reject the agent causal account on the grounds that they fail to discern any such content in their own experience. They admit that the experience of intentional agency involves a representation of oneself, but they deny that it involves any representation of the self as a primitive, substantival cause. According to such theorists, one experiences oneself as a cause only insofar as one experiences one's mental states as causally efficacious. Such a view seems to comport well with the experience of causation in general. Suppose that you see a dog cause a child to run away by barking at it. You experience the dog as causing the child to run in virtue of barking; it is the dog *qua* barker that is experienced as causing the child to run, not the dog *simpliciter*. Similarly, it seems plausible to say that it is the self *qua* bearer of certain mental states that is experienced as causally efficacious. On this view, the experience of the self as a substantival cause is parasitic on, and perhaps reducible to, the experience of event (or state) causation. Agent causationists, of course, would not be happy with this proposal, for they hold that agents are, or at least can be, primitive causes of their movements, where the force of the 'primitive' indicates that the causal relation between the agent and the movement cannot be reduced to relations between events (Chisholm; Taylor).

A second sense in which we may need to go beyond deflationism concerns the kind of object agentive experience represents the self as being (Bayne). It is sometimes suggested that in acting we experience ourselves as 'homunculi' — as 'ghosts in the machine' or 'uncontrolled controllers' (Preston and Wegner; Wegner, 'Who is the Controller'). This view is naturally paired with an agent causal account of the phenomenology of mental causation, but the two proposals should not be conflated. One could argue that although we experience ourselves as agent causes, this experience is silent about our identity conditions, and that the experience of agent causation allows that the agent to be an organism (Bishop).

I am not at all convinced that we ever experience ourselves as homunculi, strictly speaking, but I do think that there is something to the idea that in acting we experience ourselves as *things* — as substances rather than bundles. Bundle theories of the self might be correct as accounts of the self's ultimate nature, but they do not seem to have much going for them as accounts of how the self is represented in agentive experience. It's not just that the experience of the self is neutral on the question of whether or not the self is a bundle — instead, it seems to be flatly inconsistent with such a view. The experience of exerting will-power and self-control seem to be particularly problematic for such conceptions of the self (Bayne and Levy; Holton). Consider what it's like to maintain one's attentional focus,

stay awake at the wheel, or resist the temptation to have another helping of dessert. In such cases we seem to be presented to ourselves as entities that strive against some of our mental states. The content of our agentic experience appears to be at odds with those accounts of the self that would identify the self with the totality of introspectively accessible mental states. Of course, one could say that in these cases one experiences oneself as identical with a certain sub-set of our mental states — say, for example, those that we have consciously endorsed, or the like — but such a move does not have the air of phenomenal plausibility about it. Even more startling in this regard, I think, is the experience of ‘making up one’s mind’. Whatever exactly the representational content of this state is, it seems to represent oneself as a thing rather than as a bundle of events.

### 9. *The Experience of Freedom*

A third focus of recent discussion is the experience of freedom (or free will) (Nahmias, ‘Agency, Authorship, and Illusion’; Strawson, *Freedom and Belief*). Although there is clearly an intuitive sense in which we can — and often do — experience ourselves as acting freely, the satisfaction conditions of such experiences are quite obscure.

A central question here is whether experiences of freedom have compatibilist or incompatibilist satisfaction conditions. Incompatibilists regard free agency as inconsistent with determinism whereas compatibilists regard free agency as compatible with determinism. Libertarians have often assumed that experiences with incompatibilist satisfaction conditions are possible, if not altogether common. C. A. Campbell asks,

why do human beings so obstinately persist in believing that there is an indissoluble core of purely self-originated activity which even heredity and environment are powerless to affect? There can be little doubt, I think, of the answer in general terms. They do so, at bottom, because they feel certain of the existence of such activity from the immediate practical experience of themselves. (41)

More recently, Galen Strawson has claimed that there is a sense of ‘radical, absolute, buck-stopping *up-to-me-ness* in choice and action’, a sense that seems ‘indissociable from the ordinary, sane, and sober adult human sense of self’ (‘Free Agents’ 380, 394; emphasis in original; see also *Freedom and Belief*). Others who argue that some form of incompatibilism is experientially encoded include Foster, Ginet, Pink and Searle (*Rationality in Action*). Of course, such views are not universally shared. Compatibilists do not merely reject incompatibilist accounts of freedom, they typically also reject incompatibilist accounts of our *experience* of freedom.<sup>9</sup>

It is not easy to know what to make of this disagreement. Are the proponents of one (or both) sides of this debate misdescribing their experiences of freedom? Do compatibilists and incompatibilists enjoy different experiences of freedom? (If so, are these differences theory-induced, or

are they independent of theoretical commitment?) How might we resolve this debate? The difficulties that confront us in attempting to articulate what it is like to experience oneself as a free agent are not unique — other forms of experience pose similar problems — but they are particularly imposing in this domain.

One might attempt to gain some traction on these issues by asking whether it is *possible* to experience oneself as having libertarian freedom, but the question of whether it is possible for libertarian content to be experientially encoded threatens to be as intractable as the question of whether we actually experience ourselves as libertarian agents. On the one hand one might argue that we have no reason to reject the possibility of libertarian experience if we allow (as we should) that causal relations can be experientially represented. On the other hand, perhaps experiential systems are incapable of inserting a negation quite where it needs to be inserted in order to represent libertarian freedom. In order to represent an action as free in a libertarian sense one must not only represent it as undetermined by one's prior psychological properties but also as undetermined by one's physical properties — or indeed any physical properties.<sup>10</sup> And it is not obvious that experiential systems have *that* kind of representational power.

There is much of interest in the phenomenology of freedom even if we leave incompatibilism to one side. What kind of compatibilism might be consistent with, or even entailed by, the experience of freedom? Is the experience of freedom the experience of being unconstrained by external forces of a certain kind? If so, what kinds of forces? What exactly is it to experience a force as external? Compatibilist accounts of freedom are typically offered as accounts of either free will itself or of our concept of freedom, but one might well ask how well they capture the content of our *experience* of freedom.

It is sometimes suggested that there are constitutive connections between the possession of freedom and the experience of freedom (see e.g. Strawson, *Freedom and Belief*). Perhaps one must experience oneself as free in order to be free; or, perhaps, one cannot act freely if one experiences oneself as not acting freely. These are important proposals, for it is far from obvious that they are met by standard compatibilist — or, indeed, incompatibilist — accounts of freedom. Compatibilist accounts are typically advanced as analyses of what it *is* for an action to be free rather than as accounts of the *experience* of free agency, and perhaps because of this they are generally silent on how the experience of oneself as an agent might be implicated in the kinds of agentic control demanded by free agency.<sup>11</sup>

But although *prima facie* plausible, it is not clear that these proposals should be accepted. Consider the phenomenology of 'flow' that one enjoys at moments of effortless physical and creative activity. There is a case to be made for thinking that in such cases one loses any sense of oneself as a free agent; nonetheless, there is a sense in which such cases constitute paradigms of free agency. Of course, these cases are far from straightforward,

and one *might* argue that flow experience is characterised by an experience of freedom, albeit one that is recessive (that is, not the subject of attentional focus). We need a clearer conception of free agency and the experience thereof before we can pronounce with any certainty on how the two notions are related.

### *10. The Functional Role of Agentive Experiences*

What is the functional role of agentive experiences? In asking this question, I do not mean to ask about the biological function of agentive experiences — although this is certainly a question well worth asking — but about their place in our cognitive architecture. What does agentive experience contribute to the agent's cognitive economy?

Broadly speaking, there are two domains in which we might look for answers to this question. Firstly, we might look to practical reason. One proposal worth considering is that agentive experience (of a certain type) is necessary for agency (of a certain type). There may well be something to this idea, but any defence of it must reckon with the fact that at least some forms of agency appear capable of existing in the absence of agentive experience. We needn't appeal to such unusual forms of agency as that seen in the anarchic hand or utilisation behaviour to make this point — consider what it's like to navigate a busy road while engrossed in a conversation. Crossing the road involves a kind of agency, yet it is doubtful whether it is accompanied by agentive experience. But perhaps complex (should one say, 'more autonomous?') forms of agency demand certain forms of agentive experience. There is some temptation to think that high-level plans and intentions cannot be executed in the absence of the sense that one is executing them and keeping them on track. The psychologists' distinction between willed intentions and stimulus-driven intentions might be of some assistance here; perhaps willed actions must be accompanied by agentive experiences, even if stimulus-driven actions need not.

A second (and not unrelated) role for agentive experience is in theoretical reason. One might argue that first-person agentive experience is a route — perhaps the primary route — to first-person knowledge of agency. My agentive experiences not only tell me that I am trying to do something, perhaps they also tell me what I am trying to do, how I am trying to do it, and whether or not my tryings are successful. When my experiences are misleading, so too are my beliefs (unless corrected by independent information).

If this approach is on the right lines, then one might expect unusual beliefs about one's own agency to be grounded in unusual agentive experience. A number of authors have suggested that this is exactly what happens in the schizophrenic delusion of alien control (Davies et al.; Frith et al.). Patients with this delusion believe that certain of their actions are under the control of an alien force, such as God, the government, an experimenter



and so on. 'My fingers pick up the pen, but I don't control them. What they do is nothing to do with me'; 'The force moved my lips. I began to speak. The words were made for me' (qtd. in Mellors 18). Distortions of agentic experience might not provide a full account of why patients with delusions of alien control form their delusional beliefs, but it is not implausible to suppose that they will play an important role in such accounts. They also seem to support the idea that a function of agentic experience is to ground knowledge of one's own agency. Perhaps we believe that we are in control of our own bodies in large part because we experience ourselves as controlling them; when this experience is lost so too is the corresponding belief.

### 11. *The Aetiology of Agentic Experiences*

A final issue concerns the aetiology of agentic experiences. How are agentic experiences generated? This question has only recently received any attention, and even then in only a sporadic way. A few theorists have offered accounts of how particular components of the phenomenology of agency might be produced, but there has been no attempt to put together a comprehensive theory of the genesis of agentic experience.

Broadly speaking, we can identify two approaches to the aetiology of agentic self-awareness. One approach invokes a 'central' interpretive or narrative system to explain our awareness of our own agency (Carruthers, 'Illusion of Conscious Will'; Roser and Gazzaniga; Stephens and Graham; Wegner, *Illusion of Conscious Will*). Such a conception naturally falls out of a Dennettian approach to self-understanding, according to which first-person mind-reading involves taking up the intentional stance towards one's own behaviour.

An alternative approach to the aetiology of agentic self-awareness invokes dedicated sub-personal comparator systems that are concerned with motor control and production, and are largely encapsulated from one's interpretive or narrative self-conception (Haggard, 'Conscious Intention and Motor Cognition'; 'Conscious Intention and the Sense of Agency'). There is evidence that comparator systems play an important role in the generation of certain aspects of agentic experience. Sarah-Jane Blakemore and colleagues have found that schizophrenic patients with delusions of alien control are able to tickle themselves, unlike normals and schizophrenic patients with other delusions. This fact is not due to high-level differences in narrative self-conception but to low-level differences in motor-control and monitoring.

Of course, the choice between narrator-based accounts of agentic experience and comparator-based accounts is not an exclusive one, and it is entirely possible that both sorts of systems conspire to produce the full range of agentic experiences that we enjoy (Bayne and Pacherie). Exactly how the two sorts of systems might interact is very much an open question. One possibility is that one's narrative self-conception has an influence on

the high-level contents of agentive self-awareness (such as the kinds of actions that one takes oneself to be carrying out), but that the low-level contents of agentive self-awareness — what I have called the contents of agentive *experience* — are generated exclusively by low-level comparator systems. Obviously this is only one of many ways in which both low-level and high-level systems might account for agentive self-awareness. I expect that future discussions of the phenomenology of agency will be much concerned with the question of where in the architecture of cognition such states are generated.

## 12. Conclusion

My goal here has been to provide those who are approaching the phenomenology of agency for the first time with a map by means of which they might orient themselves. As will be apparent, I have been able to chart only a small portion of the terrain confronting us: vast regions of agentive self-awareness lie unmapped, and much work remains to be done before we have a clear idea of exactly what it is like to be an agent. There is good reason to do this work, for a better understanding of agentive experience promises to illuminate some of the deepest questions about both consciousness and agency.<sup>12</sup>

## Short Biography

Tim Bayne works on philosophical psychology with a particular focus on consciousness. He is completing a book on the unity of consciousness, forthcoming with Oxford University Press in 2008. Together with Axel Cleeremans and Patrick Wilken, he is editing the *Oxford Companion to Consciousness*. He has authored or co-authored papers in the *Australasian Journal of Philosophy*, *Bioethics*, *Consciousness and Cognition*, *Mind and Language*, *The Monist*, *Philosophy and Phenomenological Research*, *Philosophical Perspectives* and *Synthese*. Before taking up his current position as Lecturer at the University of Oxford and fellow of St Catherine's College he was a lecturer in the department of Philosophy at Macquarie University. He has held visiting positions at the Ecole Normale Supérieure, the Institut Jean Nicod, and the Centre for Consciousness at the Australian National University. He holds a B.A. (Hons) in Philosophy from the University of Otago and a Ph.D. in Philosophy from the University of Arizona.

## Notes

\* Correspondence address: University of Oxford and St Catherine's College, Manor Road, Oxford, OX1 3UJ, UK. Email: tim.bayne@gmail.com.

<sup>1</sup> Useful collections of papers on the topic can be found in Roessler and Eilan's *Agency and Self-Awareness*; Sebanz and Prinz's *Disorders of Volition*; Pockett et al.'s *Does Consciousness*

*Cause Behavior? An Investigation of the Nature of Volition*; Siegel's 'Symposium on the Phenomenology of Agency'. Other recent contributions to the debate not otherwise mentioned in this review include Nahmias; Proust; Siegel, 'Phenomenology of Efficacy'; White; Zhu.

<sup>2</sup> I will also restrict myself to the phenomenology of physical agency. The phenomenology of mental agency raises a number of important issues, but considerations of space force me to set them to one side here. For some thought-provoking reflections on the experience of mental agency — or rather, the lack thereof — see Strawson, 'Mental Ballistics'.

<sup>3</sup> One might argue that this line is an account of what agentive experiences is rather than a form of scepticism about agentive experience — that is, that it ought to be thought of as a form of revisionism rather than eliminativism. There is a sense in which this is true, but the fact that people might have conscious beliefs about X clearly does not show that there is a 'phenomenology of X' in any meaningful sense. If having conscious beliefs about X were sufficient for establishing a phenomenology of X, then presumably there would — or at least could — be a phenomenology of X for every X, which isn't something that we want to say.

<sup>4</sup> See the essays in Gendler and Hawthorne for a thorough examination of the representationalist approach to the phenomenal character of perception.

<sup>5</sup> Although the following paragraph contains a clear endorsement of the directive account, there are other passages in which Searle appears to leave the door open to a 'raw feels' conception of the experience of acting. In one place, Searle says that experiences of acting are intentions in action 'with certain phenomenal properties' (*Intentionality* 91f). The implication seems to be that these phenomenal properties are something over and above the intentions themselves, and hence that whatever intentional content agentive experiences have does not exhaust their phenomenal character.

<sup>6</sup> Searle's view of experiences of acting cannot be extended to account for experiences of performing an action, for an intention in action couldn't provide the intentional content of both the experience of having an intention in action and the experience of satisfying an intention in action. Might the experience of satisfying an intention in action simply be the state of performing the action in question *consciously* — that is, with certain phenomenal properties? Not if one can have an experience (as) of acting even when one is not acting.

<sup>7</sup> It is interesting to ask whether the accounts of mental causation that are advanced by, e.g. Davidson; Jackson and Pettit might also be at odds with the phenomenology of first-person agency.

<sup>8</sup> Horgan et al. use the expression 'self as source' rather than 'self as cause', but it is not completely clear to me what the distinction between cause and source might amount to.

<sup>9</sup> What about the folk? What do those without an investment in the free will debate say about their experiences of freedom? Nahmias et al. provides limited support for the claim that the folk think of their experiences of freedom in compatibilist rather than libertarian terms, but see Nichols; Turner and Nahmias for evidence that the folk also harbour libertarian inclinations about free will, if not about experiences of freedom per se.

<sup>10</sup> In his discussion of the experience of freedom, Searle (*Rationality in Action*) fails to distinguish the question of whether it is possible to experience one's actions as not having fully sufficient psychological causes from the question of whether it is possible to experience one's actions as not having fully sufficient causes of any type.

<sup>11</sup> Even those theorists who are most concerned to ground freedom in a theory of human agency, such as Fischer and Ravizza, have said very little about why the phenomenology of freedom — or, indeed, the phenomenology of agency more generally — might be necessary for free agency itself.

<sup>12</sup> Earlier versions of this paper were presented at the Research School of Social Science (The Australian National University), the Institut Jean Nicod, and the Second Cargèse Summer School on Consciousness. I am very grateful to the audiences on those occasions for their comments. I am also grateful to Brie Gertler, Jakob Hohwy, Elisabeth Pacherie and Susanna Siegel for very helpful comments on a previous draft of this paper. The research was partially funded by an Australian Research Council Discovery Grant, DP0452631.

*Works Cited*

- Bayne, T. 'Phenomenology and the Feeling of Doing: Wegner on the Conscious Will'. *Does Consciousness Cause Behavior? An Investigation of the Nature of Volition*. Eds. S. Pockett, W. P. Banks and S. Gallagher. Cambridge, MA: MIT Press, 2006. 169–86.
- and N. Levy. 'The Feeling of Doing: Deconstructing the Phenomenology of Agency'. *Disorders of Volition*. Eds. N. Sebanz and W. Prinz. Cambridge, MA: MIT Press, 2006. 49–68.
- and E. Pacherie. 'Narrators and Comparators: The Architecture of Agentive Self-Awareness'. *Synthese* 159 (2007): 475–91.
- Bishop, J. 'Agent Causation'. *Mind* 92 (1983): 61–79.
- Blakemore, S.-J., D. M. Wolpert, and C. D. Frith. 'Central Cancellation of Self-Produced Tickle Sensation'. *Nature Neuroscience* 1 (1998): 635–40.
- Bliss, J. 'Sensory Experience of Gillies de la Tourette Syndrome'. *Archives of General Psychiatry* 37 (1980): 1343–7.
- Campbell, C. A. 'In Defence of Free Will'. In *Defence of Free Will*. London: George Allen and Unwin. 1967. 35–95.
- Carruthers, P. 'Conscious Experience versus Conscious Thought'. *Consciousness: Essays from a Higher-Order Perspective*. Oxford: Oxford UP. 2005. 134–57.
- . 'The Illusion of Conscious Will'. *Synthese* 159 (2007): 197–213.
- Chisholm, R. 'Human Freedom and the Self'. *Free Will*. Ed. G. Watson. Oxford: Oxford UP, 1982. 24–35.
- Cohen, A. and J. F. Leckman. 'Sensory Phenomena Associated with Gilles de la Tourette's Syndrome'. *Journal of Clinical Psychiatry* 53 (1992): 319–23.
- Davidson, D. 'Mental Events'. Reprinted in *Essays on Actions and Events*. 2nd ed. Oxford: Oxford UP, 2001 [1970]. 207–24.
- Davies, M., M. Coltheart, R. Langdon, and N. Breen. 'Monothematic Delusions: Towards a Two-Factor Account'. *Philosophy, Psychiatry, and Psychology* 8 (2001): 133–58.
- Della Sala, S., C. Marchetti, and H. Spinnler. 'Right-Sided Anarchic (Alien) Hand: A Longitudinal Study'. *Neuropsychologia* 29.11 (1991): 1113–27.
- Estlinger, P. J., G. C. Warner, L. M. Grattan, and J. D. Easton. 'Frontal Lobe Utilization Behavior Associated with Paramedian Thalamic Infarction'. *Neurology* 41 (1991): 450–2.
- Fischer, J. M. and M. Ravizza. *Responsibility and Control*. Cambridge: Cambridge UP, 1998.
- Fodor, J. *The Immaterial Self*. London: Routledge, 1991.
- . *The Modularity of Mind*. Cambridge, MA: MIT Press, 1983.
- Frith, C. D., S.-J. Blakemore, and D. M. Wolpert. 'Explaining the Symptoms of Schizophrenia: Abnormalities in the Awareness of Action'. *Brain Research Reviews* 31 (2000): 357–63.
- Gendler, T. S. and J. Hawthorne. *Perceptual Experience*. Oxford: Oxford UP, 2006.
- Ginet, C. *On Action*. Cambridge: Cambridge UP, 1990.
- Goldberg, G. and K. K. Bloom. 'The Alien Hand Sign. Localization, Lateralization and Recovery'. *American Journal of Physical Medicine and Rehabilitation* 69 (1990): 228–38.
- Haggard, P. 'Conscious Intention and Motor Cognition'. *Trends in Cognitive Sciences* 9 (2005): 290–5.
- . 'Conscious Intention and the Sense of Agency'. *Disorders of Volition*. Eds. N. Sebanz and W. Prinz. MIT Press, 2006. 69–86.
- Hohwy, J. 'The Experience of Mental Causation'. *Behavior and Philosophy* 32 (2004): 377–400.
- Holton, R. 'How is Strength of Will Possible?'. *Weakness of Will and Practical Irrationality*. Eds. S. Stroud and C. Tappolet. Oxford: Clarendon Press, 2003. 39–67.
- Horgan, T., J. Tienson, and G. Graham. 'The Phenomenology of First-Person Agency'. *Physicalism and Mental Causation: The Metaphysics of Mind and Action*. Eds. S. Walter and H.-D. Heckmann. Exeter: Imprint Academic, 2003. 323–40.
- Hume, D. *An Enquiry Concerning Human Understanding*. Ed. Tom L. Beauchamp. Oxford: Oxford UP, 1998 [1748].
- Jackson, F. and P. Pettit. 'Functionalism and Broad Content'. *Mind* 97 (1988): 381–400.
- Korsgaard, C. *Creating the Kingdom of Ends*. Cambridge: Cambridge UP, 1996.
- Lhermitte, F. 'Utilization Behavior and its Relation to Lesions of the Frontal Lobes'. *Brain* 106 (1983): 237–55.

- Lormand, E. 'Nonphenomenal Consciousness'. *Nous* 30 (1996): 242–61.
- Marcel, A. 'The Sense of Agency: Awareness and Ownership of Action'. *Agency and Self-Awareness*. Eds. J. Roessler and N. Eilan. Oxford: Clarendon Press, 2003. 48–93.
- Mellors, C. S. 'First-Rank Symptoms of Schizophrenia'. *British Journal of Psychiatry* 117 (1970): 15–23.
- Millikan, R. G. 'Pushmi-Pullyu Representations'. Ed. J. Tomberlin. *Philosophical Perspectives IX*, 185–200. Reprinted in *Mind and Morals*. Eds. L. May and M. Friedman. Cambridge, MA: MIT Press, 1996. 145–61.
- Mossel, B. 'Action, Control and Sensations of Acting'. *Philosophical Studies* 124 (2005): 129–80.
- Nahmias, E. 'Agency, Authorship, and Illusion'. *Consciousness and Cognition* 14.4 (2005): 771–85.
- , S. Morris, T. Nadelhoffer, and J. Turner. 'The Phenomenology of Free Will'. *Journal of Consciousness Studies* 11.7–8 (2004): 162–79.
- Nichols, S. 'The Folk Psychology of Free Will: Fits and Starts'. *Mind and Language* 19.5 (2004): 473–502.
- O'Connor, T. 'Agent Causation'. *Agents, Causes, Events*. Ed. T. O'Connor. New York, NY: Oxford UP, 1995. 173–200.
- O'Shaughnessy, B. 'The Epistemology of Physical Action'. *Agency and Self-Awareness*. Eds. J. Roessler and N. Eilan. Oxford: Clarendon Press, 2003. 345–57.
- Pink, T. *Free Will: A Very Short Introduction*. Oxford: Oxford UP, 2004.
- Pitt, D. 'The Phenomenology of Cognition or What is it Like to Think that P?'. *Philosophy and Phenomenological Research* 69 (2004): 1–36.
- Preston, J. and D. M. Wegner. 'Ideal Agency: On Perceiving the Self as an Origin of Action'. *On Building, Defending, and Regulating the Self*. Eds. A. Tesser, J. Wood, and D. Stapel. Philadelphia, PA: Psychology Press, 2005. 103–25.
- Proust, J. 'Perceiving Intentions'. *Agency and Self-Awareness: Issues in Philosophy and Psychology*. Eds. J. Roessler and N. Eilan. Oxford: Clarendon Press, 2003. 48–93.
- Roser, M. E. and M. S. Gazzaniga. 'The Interpreter in Human Psychology'. *The Evolution of Primate Nervous Systems*. Eds. T. M. Preuss and J. H. Kaas. Oxford: Academic Press, 2006. 503–8.
- Searle, J. *Intentionality*. Cambridge: Cambridge UP, 1983.
- . *Rationality in Action*. Cambridge, MA: MIT Press, 2001.
- . 'The Background of Intentionality and Action'. *John Searle and His Critics*. Eds. E. Lepore and R. van Gulick. Oxford: Blackwell, 1991. 289–300.
- Siegel, S. 'The Phenomenology of Efficacy'. *Philosophical Topics* 33.1 (2005): 265–84.
- , ed. 'Symposium on the Phenomenology of Agency'. *Psyche* 13.1 (2007). <<http://psyche.cs.monash.edu.au/>>. Accessed 1 November 2007.
- Siewert, C. *The Significance of Consciousness*. Princeton, NJ: Princeton UP, 1998.
- Stephens, G. L. and G. Graham. *When Self-Consciousness Breaks*. Cambridge, MA: MIT Press, 2000.
- Strawson, G. 'Free Agents'. *Philosophical Topics* 32.1–2 (2004): 371–403.
- . *Freedom and Belief*. Oxford: Clarendon Press, 1986.
- . 'Mental Ballistics or The Involuntariness of Spontaneity'. *Proceedings of the Aristotelian Society* 103.3 (2003): 227–56.
- . *Mental Reality*. Cambridge, MA: MIT Press, 1994.
- Taylor, R. *Action and Purpose*. Englewood Cliffs, NJ: Prentice-Hall, 1966.
- Turner, J. and E. Nahmias. 'Are the Folk Agent-Causationists?' *Mind and Language* 21.5 (2006): 597–609.
- Tye, M. *Ten Problems of Consciousness*. Cambridge, MA: MIT Press, 1995.
- Wakefield, J. and H. Dreyfus. 'Intentionality and the Phenomenology of Action'. Eds. E. Lepore and R. van Gulick. *John Searle and His Critics*. Oxford: Blackwell, 1991. 259–70.
- Wegner, D. M. *The Illusion of Conscious Will*. Cambridge, MA: MIT Press, 2002.
- . 'Who is the Controller of Controlled Processes?'. *The New Unconscious*. Eds. R. Hassin, J. S. Uleman, and J. A. Bargh. New York, NY: Oxford UP, 2005. 19–36.
- White, S. L. 'Empirical Psychology, Transcendental Phenomenology, and the Self'. *Cartographies of the Mind*. Eds. S. L. White, M. Marraffa, M. De Caro, and F. Ferretti. Kluwer, 2007. 243–54.
- Zhu, J. 'Is Conscious Will an Illusion?'. *Disputatio* 16 (2004): 59–70.