

The Sense of Agency

Tim Bayne
University of Oxford and St. Catherine's College
Manor Road
Oxford OX1 3UJ
United Kingdom
tim.bayne@gmail.com

Forthcoming in F. Macpherson (ed.) *The Senses*. Oxford: Oxford University Press. This version of the paper is only a **draft**. For purposes of quotation please consult the published version.

Abstract

Where in cognitive architecture do experiences of agency lie? This chapter defends the claim that such states qualify as a species of perception. Reference to ‘the sense of agency’ should not be taken as a mere *façon de parler* but picks out a genuinely perceptual system. The chapter begins by outlining the perceptual model of agentic experience before turning to its two main rivals: the doxastic model, according to which agentic experience is really a species of belief, and the telic model, according to which agentic experience is really a species of agency. I conclude by defending the perceptual model against a number of objections to it, and by briefly exploring its implications for the question of how to approach the study of perception.

The motion of our body follows upon the command of our will.

Of this we are at every moment conscious.

David Hume, *An Enquiry Concerning Human Understanding*, 1748

1. Introduction

Hume was perhaps guilty of a certain degree of hyperbole in claiming that we are continually aware of ourselves as agents, but it is certainly true that agentic self-awareness is a frequent component of the stream of consciousness. The question with which I am concerned in this chapter is where we ought to locate states of agentic self-awareness—‘experiences of agency’, as I will also call them—in cognitive architecture.

To a first approximation, we can distinguish between three broad departments of the mind: perception, thought and action. Perception processes information about the world (where this includes the state of the subject’s own body) and makes it available for cognitive consumption. Perceptual systems include among their number not only the five traditional senses but also those systems responsible for the various kinds of bodily sensations, such as experiences of joint position, pain, nausea, hunger, thirst and the need for oxygen. The department of thought is responsible for doing something with the information that it receives from perception—it is in the business of theoretical and practical reasoning. And the function of the department of agency is to implement the commands that it receives from thought. Its job is to plan and execute action in conjunction with the department of thought.

Where in this picture does agentic experience fall? I will argue that agentic experiences are best thought of in perceptual terms—they are the products of a dedicated perceptual system. Just as we have sensory systems that function to inform us about the distribution of objects in our immediate environment, damage to our limbs, and our need for food, so too we have a sensory system (or systems) whose function it is to inform us about facets of our own agency.

The chapter will unfold as follows. The following section introduces the notion of agentic self-awareness and presents a sketch of the perceptual model; the remainder of the chapter constitutes an argument for the model, chiefly in the form of showing that its central rivals—the doxastic and telic models—are untenable. Section three examines the doxastic model of agentic self-awareness, according to which agentic self-awareness is solely a matter of judgment and is to be located within cognition. Section four examines the telic model of agentic self-awareness, according to which experiences of agency are

to be located within the action system. Neither the doxastic nor the telic model provides us with a plausible analysis of agentic experience. Section five reinforces the intuitive case for the perceptual model of agentic experience by showing that a number of the most important objections to it can be met. My overall aim in this chapter is not to present a knockdown case for the perceptual model, but rather to show that it ought to be regarded as a live option. On this view, talk of ‘the sense of agency’ is no mere *façon de parler* but picks out a genuine sensory modality, at least in a fairly broad sense of the term.¹ Our conception of the senses ought to be wide enough to include room for the idea that there might be a system (or systems) that is in the business of representing one’s own agency.

2. The Perceptual Model of Agentic Experience

Consider the following vignette:

It’s your first day as a waiter/waitress, and you are pouring water into a glass from a jug. As you pour the water, you experience yourself as an agent. You experience yourself as someone who is doing something, rather than someone to whom things are merely happening. The experience of yourself as an agent is modulated in various ways. In pouring the water, you experience yourself as trying to avoid being distracted by the commotion in the kitchen. This attempt to control your attention is experienced as effortful—as difficult. You also experience yourself as having some degree of freedom or autonomy in controlling your attention. Having poured the water, you deliberate as to whether or not you should ask your patrons if they are ready to have their order taken. This act of deliberation has a distinctive experiential character; it too colours what it is like to be you during this episode.

This vignette contains reference to multiple forms of agentic self-awareness—states that I will refer to as ‘agentic experiences’.² There is no agreed taxonomy of such states.

¹ The ‘sense of agency’ is sometimes used as a superordinate term for agentic experiences in general (e.g. Marcel 2003) and as a label for a particular component of agentic experience—roughly, the experience of authoring an action (see e.g. Pacherie 2008). As should be clear from the foregoing, I use the term for a faculty (sensory system) that produces agentic experiences rather than as a term for (any species of) the experiences produced by this system.

² ‘The phenomenology of first-person agency’ is also used as a synonym for ‘agentic experience’. I prefer the former label over the latter, for ‘phenomenology’ has often been used to refer to a method for studying consciousness rather than as a term for the phenomenal character of experience.

Among the various terms in use are the following: 'the experience of control', the 'experience of free will', 'the experience of acting', 'the experience of volition', 'the experience of conscious will', 'the experience of efficacy', 'the experience of mental causation', 'the experience of effort', and 'the feeling of doing'. Some of these terms are best understood as synonyms of each other; others pick out distinct elements or aspects of agentic experience. What they all have in common is an attempt to demarcate forms of experience that can be individuated by reference to features of one's own agency.

I take agentic experience to have as its core the experience of a particular movement or mental event as realizing one's own agency. Surrounding this core are a number of other experiential components, the precise nature and number of which is very much an open question. Actions are sometimes experienced as effortful and at other times as effortless. Actions are sometimes experienced as deliberate and at other times as spontaneous. More controversially, we often experience our actions as free and autonomous—as 'up to us' in some way. There is debate about the content of such experiences—are they compatibilist or incompatibilist in nature?—but a strong case can be made for thinking that some form of free will or autonomy is experientially encoded. One might add to this list the phenomenology of deliberating, and perhaps also the phenomenology of deciding.³

The existence of agentic self-awareness can be highlighted by drawing attention to pathologies of agency in which it is lost or disturbed. In the anarchic hand syndrome, patients will complain that they have lost control of a hand; that will report that one of their hands has acquired a "will of its own" (Banks *et al.* 1989; Della Sala *et al.* 1991; Gasquoine 1993; Goldberg & Bloom 1990; Pacherie 2007). The hand in question might also appear to act in an anarchic fashion, by (say) grabbing food from a stranger's plate at a restaurant. Patients will deny that such actions are theirs, and may take steps to prevent the hand from moving by sitting on it or tying it down. Detailed reports of what it is like to have an anarchic hand are rare, but it is plausible to suppose that the syndrome involves an experience of alienated agency—an experience of the hand as acting on its own accord.

The schizophrenic delusions of thought-insertion and alien control also involve disturbances of agentic experience (Mellors 1970; Frith *et al.* 2000). Patients suffering from these delusions complain that some of their thoughts and actions are no longer their own, or are at least no longer under their control. "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

³ See Bayne (2006), Bayne (2008), Bayne and Levy (2006), Hohwy (2004), Holton (2006), Horgan *et al.* (2003), Marcel (2003), Nahmias *et al.* (2004), Pacherie (2008) and Siegel (2005) for further discussion of the contents of agentic experience.

to speak. The words were made for me” (Mellors 1970: 18). “My grandfather hypnotized me and now he moves my foot up and down.” “They inserted a computer in my brain. It makes me turn to the left or right.” (Frith et al 2000: 358). It is highly plausible to suppose that these delusions are at least partially grounded in abnormal experiences of agency (Pacherie et al 2006). And if these syndromes involve experiences of alienated agency, then there is some reason to suppose that unimpaired agency might be accompanied by experiences of intact agency. Of course, the fact that there are experiences of alienated agency does not entail that there are—or even could be—experience of (un-alienated) agency, but it surely provides some support for that thesis.

Where in the cognitive architecture should we locate such experience? They are not located within the central cognition, nor are they located within the systems responsible for programming and executing actions, nor are they located within the high-level reaches of any of the standard perceptual modalities (although they are intimately connected to the other modalities). Instead, such states are the products of a dedicated perceptual system (or systems). These systems involve forward models of action control, and operate along the following lines (Frith et al., 2000a, 2000b; Blakemore and Frith, 2003). The forward models are fed a copy of the agent’s motor commands, and are then put to the following uses. Firstly, they are used to predict the sensory consequences of the agent’s movement. These predictions are used to filter sensory information and to attenuate the component that is due to self-movement. Secondly, the copy of the motor commands is also sent to a comparator (or comparators), so-called because they compare the predicted sensory consequences of the movement with sensory feedback. When there is a match between predicted and actual state, the comparator sends a signal to the effect that the sensory changes are self-generated; when there is no match (or an insufficiently robust match), sensory changes are coded as externally caused. Crucial to the comparator approach is the notion that agentic awareness can be generated by mechanisms that need not—and typically will not—have access to fully-fledged intentions. The comparator-based story may not account for each of the components of agentic experience, but there is good reason to think that it lies at the heart of the phenomenon.⁴

Where in the cognitive architecture do these comparators fall? They lie between the standard perceptual systems and the motor system, for they take both perceptual representations and motor representations as input. Most importantly, the states that they generate have the functional role of perceptions. Like other perceptual systems, the

⁴ For further details see Bayne & Pacherie (2007), David et al (2008), Haggard (2005), Hallett (2007) and Pacherie (2008).

function of the sense of agency is to generate representations of some domain—in this case, one’s own agency—and make those representations available to the agent’s cognitive systems in an experiential format.

It is important to note that the perceptual model of agentic experience is *not* committed to a perceptual account of introspection. The perceptual model is a model of agentic experiences themselves, it is not a model of our access to agentic experiences. Giving an account of agentic experience is one thing, giving an account of introspection is quite another. Perceptual approaches to introspection are decidedly unpopular, but that unpopularity should in no way be taken to discredit the perceptual model of agentic experience. It is equally important to note that the perceptual model is not intended to cover everything that might be subsumed by the term “agentic self-awareness”. There is a broad sense of agentic self-awareness that includes the conceptually articulated awareness of what one is trying to do (tie one’s shoelaces, look for a book, reach for the salt). Although a full account of agentic self-awareness needs to address high-level, conceptually articulated forms of agentic awareness of this kind, I will not take on that task here. The perceptual model is intended to apply only to fairly low-level forms of agentic awareness of the kind introduced above.⁵

That, in a nutshell, is the perceptual model. I will flesh the model out in more detail below (section five), but I turn now to its rivals.

3. The Doxastic Model of Agentic Experience

There will be those who think that the perceptual model gets off on the wrong foot by thinking of agentic self-awareness in experiential terms. According to what we might think of as the doxastic model of agentic self-awareness, such states are not really forms of experience at all; instead, they are states of thought. What it is to be aware of oneself as an agent is simply to have a conscious thought to that effect.

The doxastic model strikes me as intuitively implausible. Horgan *et al* (2003: 324) describe agentic experience as “an aspect of sensory-perceptual experience, broadly construed”, and that seems right. The experience of acting is not exhausted by the conscious judgments one might have about what or how one is acting, if indeed one has any conscious judgments of that nature at all. Acting may frequently involve conscious judgment about what one is doing, but agentic self-awareness is not primarily a matter of judgment. Rather, it has the transparency, immediacy and directness that characterizes

⁵ See Bayne and Pacherie (2007) for some discussion of how low-level and high-level agentic self-awareness might be related.

our sensory engagement with the world. There might be good reasons to reject this view, but any such rejection would come at a significant cost to intuition.

Proponents of the doxastic model who are unmoved by the foregoing might find the following more convincing. If the doxastic view were right, then pathologies of agentive self-awareness would be pathologies of judgment rather than of experience. We should hold that the anarchic hand patient syndrome involves only a departure in the agentive judgments that the patient is prepared to make. Whereas a person whose hand has just taken food from a stranger's plate would normally be prepared to judge that they had performed the action (and perhaps confabulate a rationalization for it), the anarchic hand patient is not prepared to make this judgment and insists that she did not move her hand—it moved of its own accord.

But this doesn't seem to capture the situation. Although it is certainly true that the anarchic hand patient isn't (typically) prepared to admit that the action was hers, it seems plausible to appeal to agentive experience—or the lack thereof—in order to *explain* why she denies having authored the anarchic actions. It is the fact that the normal and expected experience of *doing* has been replaced by an experience of *happening* which drives the patient to judge that the action is not hers. Whether or not the patient's denial of authorship over the action is erroneous—and there is room for debate about this—the patient's pathology is primarily one of experience rather than judgment.

This objection can be reinforced by noting that experiences of alienated agency are unlikely to be cognitively penetrable. Suppose that the anarchic hand patient *does* take herself to be the agent of her anarchic actions. She might reason to herself as follows: "The movements of my anarchic hand are not guided by anyone else. They are actions, and where they are actions there must be an agent. So, these actions must be mine." Will forming the belief that she is the author of her 'anarchic' actions suffice to correct her agentive experience? That seems highly implausible. The patient can no more restore the missing experience of agency by forming the belief that her anarchic actions are truly her own than you or I can correct our visual experiences of the Müller-Lyer illusion by forming the belief that the lines in question are equal in length. Cognitive impenetrability, one of the key markers of the perception-cognitive boundary, tells firmly against the doxastic analysis of agentive self-awareness.

The point generalizes. Patients suffering from depersonalisation seem also to have lost the normal experience of themselves as agents. They "complain that they no longer have an ego, but are mechanisms, automatons, puppets; that what they do seems not done by them but happens mechanically" (Schilder 1953; Sierra & Berrios 1998). Nonetheless, they

suffer from no delusions concerning agency: they know that their actions are their own. Phantom limb patients display what is in some sense the converse dissociation. They know that their 'limb' is but a phantom, but they continue to experience themselves as acting with it. The history of philosophy provides us with additional examples of the cognitive impenetrability of agentic experience. Malebranche believed that God was the direct cause of his movements, but his commitment to occasionalism surely did not rob him of the experience of being the author of his deeds. Hard incompatibilists reject the reality of free will, but this has little impact on their experiences of freedom (see e.g. Strawson 1986). Perhaps judgment can penetrate agentic experience under some conditions, but it seems clear that on the whole agentic experience exhibits the kind of doxastic impenetrability that is characteristic of perception. Agentic experiences are as distinct from the judgments to which they give rise as are visual experiences from the judgements to which they give rise. Experience is an "invitation to belief", but one should not confuse an invitation itself with that to which it is an invitation.

In response to this, the proponent of the doxastic model might hold that even if agentic experience is not a species of judgment strictly speaking, it does fall on the cognitive side of the perception-cognition divide. But what kind of cognitive state might agentic experience be if not a judgment? The only answer I know of is that it might be an entertaining of some kind. Suppose, when asked what the capital of Ethiopia is, the name "Addis Ababa" suddenly occurs to you. Do you judge that Addis Ababa is the capital of Ethiopia? Perhaps not. You might have no commitment whatsoever to the truth of that thought, it's merely that you 'found yourself' with it when the question was put to you. Might agentic self-awareness involve states of this kind—thoughts that simply occur to one?

I think not. The difference between the case just discussed and agentic experience is as sharp as the difference between daydreaming about what it would be like to fly through the air unaided and actually experiencing oneself as being in this state. When it comes to daydreams and flights of fancy, one has no tendency to take oneself as being in touch with how things really are; perception, by contrast, does involve such a sense. And that is surely what agentic experience is like. It is committal. The experience of raising one's arm does not leave it open as to whether one might really be raising one's arm but has the force of an assertion: "I am raising my arm".

4. The Telic Model of Agentic Experience

The telic model of agentic experience grounds such states within the motor system. Proponents of this model might agree with Horgan et al's description of agentic

experience as “an aspect of sensory-perceptual experience, broadly construed”, but they will insist that this description must be construed very broadly indeed if it is to be accurate, for on this account the structure of agentive experience is diametrically opposed to that of perceptual states.

Following Anscombe (1957) and Searle (1983), let us distinguish two ways in which mental states can be satisfied. *Thetic* states are satisfied when they fit the world; they have a mind-to-world direction of fit. Judgments are paradigm of thetic states, for they are true (or veridical) when they fit the world and otherwise false (or non-veridical). Thetic states are in the business of responding to changes in the world. *Telic* states are satisfied when the world fits them; they have a world to mind direction of fit. Desires are the paradigms of telic states. They are in the business of bringing about a certain state of affairs. They are satisfied when they succeed; otherwise they remain frustrated.

The perceptual model takes agentive experiences to have a thetic structure: they are representations (or ‘presentations’) that matters are thus-and-so. By contrast, the telic model takes such states to have a telic structure. Searle himself endorses the telic model.

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 Intentional 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. (1983: 88, emphasis in original; see also 1983: 123f.)

Although Searle is perhaps the most explicit advocate of the telic line, a number of other authors appear to have some sympathy for at least a limited version of the view (O'Shaughnessy 2003; Peacocke 2003).⁶

There are multiple ways in which we might get a fix on the dispute between the thetic and telic analyses of agentic experience. It can be read as a dispute about how to read the 'of' in the phrase 'the phenomenology *of* agency', a phrase that is often used to refer to agentic experience. The thetic theorist reads this 'of' intentionally: the phenomenology of agency involves experiences that are intentionally directed towards agency. By contrast, the telic theorist reads this 'of' possessively: the phenomenology of agency is a matter of one's actions (or tryings) enjoying experiential character. The dispute can also be captured by appeal to the contrast between transitive and intransitive conceptions of experience. The thetic theorist takes a transitive view of agentic experience (the phenomenology of agency is a matter of experiencing that certain things are the case), whereas the telic theorist takes an intransitive view of it (the phenomenology of agency is a matter of one's actions, intentions and the like having phenomenal properties).⁷

How might we determine whether agentic experience is thetic or telic? It is unclear whether introspection is competent to pronounce on this question. According to one widely endorsed view, introspection provides us with access only to the contents of experience. Since the debate between the thetic and telic accounts concerns the structure rather than the contents of experience, introspection may be impotent here. Of course, this view of introspection is highly contested, and many would argue that there are introspectively accessible differences between experiences that outrun their contents. (What it is like to see something as square is different from what it is like to feel it as square, but arguably the two states share the same content.) Whether or not the current

⁶ I find Peacocke's (2003, 2008) position difficult to pin down. On the one hand, Peacocke denies that agentic experience (what he calls "action-awareness") can be understood in perceptual terms. At the same time he grants that agentic experience is representational "...in the sense that in enjoying action-awareness, it seems to the subject that the world is a certain way" (2008: 247). I am unsure how to reconcile these two elements of Peacocke's account, and it may be that his view is not really a form of the telic model despite appearances to that effect.

⁷ We might note in passing that the telic model of agentic experience is at odds with standard representationalist conceptions of phenomenal character, for on such accounts the phenomenal character of a mental state is fixed by how it represents the world as being (see e.g. Dretske 1995; Tye 1995). However, the telic model is consistent with the spirit of such analyses for it allows that facts about intentional content might fix facts about phenomenal character.

debate falls within the domain in which introspection *could* deliver a verdict, it is clear that introspection has not proven particularly useful in resolving it. At best introspection is silent on the question of whether agentive experience has thetic or telic structure, at worst it inclines theorists in a telic direction and others in a thetic one.

We might hope to get more traction on this debate by asking how agentive experiences could fail. States with a mind-to-world direction of fit fail by being frustrated, whereas states with a world-to-mind direction of fit fail by misrepresenting how things are. In what way—or ways—do agentive experiences fail?

At least some forms of agentive experience fail (when they do) by misrepresenting. This is surely what we ought to say about experiences of free will. Whatever the exact content of experiences of free will, it is surely an open question whether we are free in the way(s) in which we experience ourselves as being free. Searle himself takes it to be very much an open question whether we actually possess the libertarian freedom that he thinks we experience ourselves as having (Searle 2001; see also Ginet 1997). Doubts about the veridicality of experiences of free will are most pressing for those attracted to an incompatibilist analysis of their content, but they can also be raised within the context of compatibilist analyses.

Experiences of effort can also be non-veridical: in principle, it seems possible that one might experience effortful tasks as effortless and vice-versa. In fact, this possibility might not be merely theoretical. Naccache *et al.* (2005) describe a patient who performed normally on measures of cognitive load but experienced tasks involving high cognitive load as no more effortful than tasks involving low cognitive load. It is natural to suppose that in experiencing these tasks as effortless this patient was misrepresenting her own cognitive processing—they were effortful, despite being experienced as effortless.

But what about the core components of agentive experience? Consider William James's case of a 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 he is surprised to discover that his arm has not moved. Although he tried to raise his arm, this attempt was a failure. What did James's patient *experience*? Searle describes his experience as an experience of '*trying but failing to raise one's hand*' (Searle 1983: 89; emphasis in original), but that seems wrong. Arguably, the patient was surprised to discover that his hand had not gone up because he experienced himself as raising it—the patient had an experience *as* of raising his arm. If this is right, then the experience of acting cannot be identified with acting '*accompanied by certain phenomenal properties*' (Searle 1983: 92), for in this case there was no action for phenomenal properties to accompany. The case

appears to demand a thetic analysis. On discovering that his arm has not moved, James' patient ought to treat his experience of acting in just the way he treats his perceptual experience of illusions—*viz.*, as misrepresenting some aspect of reality.

The proponent of the telic model might reply that the patient's experience misleads him only insofar as it represents the *movement* of his arm—and there is nothing agentive, as such, about that content. We can see that this response fails by considering the following variant of James's case. Suppose that the patient is paralyzed; he can try to raise his arm, but his trying will not be causally efficacious. However, a mischievous scientist with brain-reading technology can detect the patient's tryings, and intervenes in his motor cortex so as to ensure that his tryings are successful. (This, of course, is roughly how the occasionalists thought all agency worked.) The patient will experience himself as raising his arm—after all, his arm goes up as and when he tries to raise it—but there is surely something non-veridical in his experience. This non-veridical element cannot be identified with the representation of the trajectory or location of his arm, for in this case these elements are veridical. Rather, it must concern his experience of *agency*—it must have something to do with the agent's role in producing the movement.⁸

Counter-examples aside, there is something rather odd with the thought that experiences (as) of acting successfully (unsuccessfully) could be identified with successful (failed) tryings. Why should the agent have any first-person access to the success of her tryings, merely in virtue of having successful tryings? Consider, as a parallel, the proposal that we identify the experience of one's perceptions as veridical (non-veridical) with the veridicality (non-veridicality) of one's perceptions. Such a proposal is clearly a non-starter: the veridicality of a perception is one thing, the experience of it as (non)-veridical is another. Identifying the experience of one's intentions as satisfied or frustrated with the event of them being satisfied or frustrated is similarly implausible.⁹

⁸ One can also arrive at this conclusion via a highly restrictive conception of actions, according to which actions are identified with tryings (Hornsby 1980). On this view, James' patient does in fact raise his arm, it's just that the world fails to cooperate with his action. I will assume here that this account of actions is incorrect, at least insofar as it is taken as an analysis of the content of the experience of acting.

⁹ The telic theorist could go disjunctive at this point. She could reject the 'common factor' assumption, according to which agentive experiences belong to a single kind whether or not they occur in the context of successful intentions ('the good case') or unsuccessful intentions ('the bad case'). Wading into the murky waters of disjunctivism is beyond the scope of this chapter, but suffice it to say that if the telic theorist needs to go disjunctivist then I do not want to go telic.

At this point, the proponent of the telic model might restrict the model to experiences of trying. Perhaps the experience of trying to do something is not a state that represents oneself as trying to do something, but simply a trying 'that possesses phenomenal character'. In fact, there is some reason for taking Searle's phrase 'the experience of acting' as referring to the experience of trying to act as opposed to the experience of acting. At any rate, Searle does appear to think that experiences of trying to act have telic structure.

As I read him, Searle provides two arguments for this proposal. Firstly, he claims that it is part of the content of experiences of trying that they represent themselves as bringing about those actions with which they are associated: "If it [the experience of acting] 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" (1983: 88; see also 124). If this claim were true it would provide an attractive argument for the telic account, for as we saw in the case of the mischievous scientist, the satisfaction conditions of trying to raise one's arm appear to include a self-referential component. Arguably, it is part of the content of the trying that it itself plays a direct role in causing one's arm to go up. It doesn't seem implausible to identify experiences of trying with tryings if they share the same self-referential content.

But do they share the same self-referential content? I doubt it. I can discern nothing in the phenomenology of the experience of trying to φ which represents that very experience as enabling me to φ . I am not convinced that experiences of trying contain any self-referentially causal content, but if they do I am inclined to think that they represent themselves as being caused by tryings, rather than representing themselves as causing actions. At any rate, I do not share Searle's intuition that in raising my arm I experience that experience as itself causing my arm to rise. As far as I can tell, the only thing I experience as causing my arm to rise is me.

Searle's second argument for identifying experiences of tryings with tryings involves an appeal to failures of agency. He writes: "when we have an experience [of acting] 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." (Searle 1983: 88). As I read it, the argument is that the telic account must be right because it correctly predicts that there is only a single "joint of intentionality" between experiences (as) of trying and actions. By contrast, the thetic account posits two such joints: one between experiences of trying and tryings themselves, and another between tryings and their realization. If the thetic account were correct, it would follow that an agent could experience herself as trying to φ even when she did not φ for one of two reasons: either her attempt at φ -ing was unsuccessful, or her experience as of attempting to φ was non-veridical. Searle's argument seems to be that because we

never countenance this second possibility it follows that experiences of tryings must be telic.

Searle might be right to suggest that we do not generally recognize a gap between experiences of tryings and trying themselves, and in this sense he may be right to suggest that our folk conception of experiences of trying is, in some sense, telic. But perhaps the folk model of experiences of trying is incorrect. Perhaps we should be open to the idea that experiences of trying—no less than experiences of actually acting—can lead us astray.¹⁰ Consider anosognosia for hemiplegia, a condition in which one is unaware that one is paralyzed on one side of one's body. When requested to perform an action that involves the paralyzed limb (such as clapping one's hands together) patients will typically insist that they have performed the requested action, despite the fact that they have moved only one hand. Among the various accounts of anosognosia for hemiplegia is the feedforward account, according to which anosognosia for hemiplegia involves a deficit in the motor effector system (Heilman 1991, Gold *et al.* 1994; Adair *et al.* 1997). Proponents of this account argue that damage to the motor effector system has resulted in the loss of the ability to form intentions to move; in other words, patients no longer try to move. Nonetheless, it seems plausible to suppose that these patients retain the experience of trying to move their limb, for they insist that they are moving (or have just moved) the limb. In short, there is some reason to think that anosognosia for hemiplegia may involve a situation in which a person's experiences of trying to do something are non-veridical: there is no trying here, not even a frustrated one, only an experience thereof. This account of anosognosia for hemiplegia might not be correct but it does seem to be *coherent*, and its coherence suffices to undermine the telic account of experiences of trying.

A full analysis of the telic model goes beyond the scope of this chapter, and I have not argued that agentic experience could not—or indeed does not—have telic components. Rather, my goal has been the rather more modest one of defending a thetic analysis of an important range of agentic experiences—most centrally, the experiences of acting and of trying to act. At most, the telic model is true of only a very limited range of agentic experiences.

¹⁰ Perhaps the representational nature of experiences of trying has been overlooked for the same reason that the representational nature of experiences of pain has been overlooked: in each case, having the experience of being in the target state provides by far the best evidence that we are in the target state.

5. Objections and Replies

I conclude by addressing three objections to the perceptual model. The first objection concerns the relationship between agentive experiences and their intentional objects. The intentional object of a perceptual experience is typically independent of the experience itself. My visual experience of a dog is one thing, the dog itself is quite another; my auditory experience of a tree falling in the forest is one thing; the falling of the tree is quite another. One might take this independence—a fact that is arguably represented in the very contents of perception—to be a necessary feature of perceptual experience.

Are agentive experiences independent of their intentional objects? Some say “no”. Mossel has recently rejected the perceptual model of what he calls “sensations of acting” on the grounds that such states are parts of actions, and hence are not independent of actions in the way in which they would need to be were they to qualify as perceptions of agency (Mossel 2005: 134). Mossel is not alone in suggesting that agentive experiences are in some way internal to actions themselves. Searle states that “...if my arm goes up, but goes up [without the experience of acting], I didn’t raise my arm, it just went up” (Searle 1983: 88, see also p. 95). Ginet strikes a similar note, claiming that a simple mental event is an action “if and only if has an intrinsic phenomenal quality”, which he dubs “the actish quality” (1997: 89).¹¹

There is *something* to this line of argument. Many people are reluctant to regard physical or mental happenings that are unaccompanied by a basic ‘experience of doing’ as actions. This reluctance is manifest in our intuitive response to the anarchic hand syndrome, for as we saw, we have some inclination to deny that anarchic actions really do qualify as the actions of the patient in question. Arguably this intuitive resistance is best explained by supposing that agentive experience of some kind is essential to agency.

But there is another side to the story. If anarchic hand actions do not qualify as the patient’s actions whose actions are they? One might deny that they are actions at all, but that seems implausible. Anarchic hand movements certainly *look* like actions, for they can be evaluated for success. (The attempt to take food from a stranger’s plate might be frustrated by the intervention of a waiter!) Peacocke (2003) has suggested that anarchic-hand actions are ‘orphans’—actions without an agent. This position is also implausible, for we tend to think that actions must come attached to agents. So perhaps we should admit that anarchic hand actions belong to the patient despite the fact that she or he

¹¹ Searle does not seem to have a settled position on this issue, for certain passages in *Intentionality* appear to allow that actions need not be accompanied by the experience of acting (see e.g. 1983: 92).

experiences no sense of agency towards them. Some support for the view that actions need not be accompanied by agentive experiences derives from what Bach (1978) calls 'minimal actions', such as tapping absentmindedly whilst listening to a paper, or walking to the shops whilst one's mind is on other matters. It is an open question whether minimal actions always accompanied by an experience of agency. Although I myself am inclined to think that they are generally accompanied by at least a recessive or background kind of agentive experience, there are those who hold that minimal actions are typically executed without any agentive experience whatsoever. Would such actions fail to qualify as bona fide actions? I do not see why not.¹²

Although I reject the independence objection, I do not deny that agentive experiences of various stripes are often causally implicated in agency. Clearly, the experience of losing control of one's actions can lead one to attend more closely to what one is doing, which may in turn affect how (or indeed whether) one continues to act. The experience of a certain action as effortful might lead one to either exert more effort or, alternatively, to abandon doing what one is trying to do. And perhaps the experience of acting freely or autonomously is in some way implicated in what it is to act freely or autonomously. Just how agentive experience might be implicated in agency itself is still quite obscure, but there is good reason to think that much ordinary agency would be unrecognisable in the absence of agentive experience; indeed, perhaps we would not even have the tools to conceive of ourselves as agents in the absence of agentive experience. Granting all this, there remains enough independence between agentive experiences and their conditions of satisfaction for it to be possible, in principle if not in practice, that such states are non-veridical.¹³

¹² Although I have argued that the states of affairs represented by agentive experiences are independent of them, it may not be possible to give an analysis of what an action is that does not appeal in some way to agentive experience. But even if this should turn out to be the case, it would still be possible for agentive experiences to be non-veridical. As a parallel, consider colours. Response-dependence conceptions of colour hold that there is no account of what colours are that doesn't ultimately appeal to colour experiences. Nonetheless, such accounts allow that it is possible for an object to appear to be green without being green. Similarly, a response-dependence conception of actions will allow that it is possible for something to appear to be an action without being an agent (and vice-versa).

¹³ Perhaps the folk notion of action is not determine enough for the question just examined to have an unequivocal answer. Intentional goal-directed movement is typically accompanied by agentive experience, and our commonsense notion of action might simply having nothing to say about

A second objection to the perceptual model focuses on the question of whether agentive experience involves the operation of a perceptual organ of the requisite kind. On the face of things, there does not seem to be anything that stands to agentive experience as the eyes stand to seeing, the ears stand to hearing, or the skin stands to touch. Might this fact scuttle the perceptual model?

To address the question we need to consider how the notion of a perceptual organ is to be understood. According to O'Dea's very interesting analysis (this volume), a perceptual organ is a mechanism over whose operation one has some degree of intentional control and whose mode of operation is implicitly represented in the very experiences it produces. On this proposal, perceptual experience involves some form of self-intimation of its own mode of acquisition: it is part of the content of visual experience that it is acquired via the deployment of the eyes, it is part of the content of tactile experience that it is acquired via through the skin, and so on. O'Dea's account builds on a Gibsonian conception of perception, according to which perceptual organs are not mechanisms that passively register their input but tools by means of which we explore our environments and track objects.

O'Dea's analysis captures an important feature of the traditional senses—a feature that sets “the famous five” apart from other forms of sensory experience—but it is not the only viable notion of a perceptual organ. Consider proprioceptive, vestibular and nociceptive experience. The mechanisms that lie behind these forms of experience are not open to view or amenable to intentional control. There is nothing that one can do in order to get a better perspective on one's vestibular environment, and there is no straightforward sense in which one can track the state of one's viscera. Nonetheless, there is a respectable sense in which these forms of experience do qualify as perceptual, for they are ways of gaining information about the state of one's environment (in this case, the state of one's own body). Furthermore, it is not clear that O'Dea's analysis even applies to each of the famous five without qualification. Perhaps animals and neonates can enjoy perceptual experience—they can see things, touch things, and so on—without any awareness of how they come by this information. O'Dea's analysis does capture something important, but we should resist the suggestion that it gives us the only legitimate analysis of what it is to be a perceptual organ.

One could take a perceptual organ to be a dedicated mechanism that takes as input energy of some kind and generates representations in an appropriate format, at least some of

whether intentional goal-directed movement that fails to be accompanied by agentive experience qualifies as an action.

which are experiential. But do the mechanisms responsible for agentic experiences do even that? On the comparator-based account introduced earlier, the sense of agency takes as input not raw energy but motor representations and representations drawn from the other perceptual modalities. So, there is a real difference between the mechanisms responsible for agentic experiences and those that generate (say) proprioceptive, vestibular and nociceptive experiences: the latter involve sensory transducers whereas the former do not. Nonetheless, I am not inclined to regard this difference as particularly significant in the present context. The appropriate response to it is to draw a distinction between basic perceptual systems and non-basic perceptual systems, where the former take as input forms of energy and the latter take representations as input. The five traditional senses (and various forms of proprioception and interoception) would qualify as basic perceptual systems, whereas the sense of agency would qualify as a non-basic perceptual system. But a non-basic perceptual system is nonetheless still genuinely perceptual.

A third objection to the perceptual model is perhaps the hardest to pin-down with any precision. Put most generally, the worry is that any perceptual approach to agentic experience will inevitably drain such experiences of their agentic nature. As Korsgaard puts it, “to experience something is (in part) to be passively receptive to it, and therefore we cannot have experiences of activity as such” (1996: 204). More enigmatically, O’Shaughnessy writes:

If one is to relate as observer to anything then one has to be ‘without it’, whereas if one is intentionally to do anything then one has to be ‘within’ the action we are attempting to observe, in which case we have an entirely empty and self-delusive experience of observation....or else we remain ‘without’ in some less serious sense and genuinely seem to observe the action. But remaining ‘without’, we lose the action as ours in gaining the observation: we lose any ‘withinness’ (O’Shaughnessy 1980: 31).

I must admit that I fail to feel the force of the objection. There is no doubt some sense in which perception is passive, but it is surely a gross *non sequitur* to suppose that because a state is passive it cannot represent agency. After all, judgement is also passive, but we can clearly form judgments about our own agency. If it is possible for one’s own agency to be represented in thought why should it not also be possible for it to be represented in experience? Why exactly would the perceptual representation of agency undermine its ‘withinness’?

I do not have a full answer to this question, but I suspect that the worry that lies behind these claims concerns the thought that on the perceptual model it must always be an open question whether the action being perceived as one's own really is one's own. Perceptual experience is not immune to error relative to the first-person. For example, one can raise the question of whether the legs that appear to one in visual experience are one's own or those of another (Shoemaker 1968). But, so the thought goes, agentic experience *is* immune to error relative to the first-person. The question of whether the actions of which one is aware in experience are one's own or those of someone else is puzzling at best and downright incoherent at worst. So, agentic experience is not perceptual. Engaging with this worry would lead us into deep waters that I cannot even begin to chart in this chapter, suffice it to say that the distinction between logical and de facto immunity to error might of some help to us here. The objection requires that agentic experience be logically immune to error, but it may in fact be only de facto immune to error, in the way in which proprioception arguably is. No doubt there are further objections to the perceptual model, but I leave to others the task presenting them.

6. Conclusion: The Sense of Agency and the Famous Five

I have argued that there is much to be said on behalf of a perceptual model of agentic experience. The model does justice to its phenomenology; it receives some support from cognitive science; and is able to meet the most pressing objections to it. Assuming that the perceptual model is on the right track, what lessons might we draw from it?

There are lessons to be drawn in two domains. Firstly, the perceptual model opens up questions that are obscured by other analyses of agentic experience. How reliable is the sense of agency? Does agentic experience represent agency as it really is, or does the sense of agency lead us into error—not just occasionally but systematically? To what degree is the manifest image of agency as encoded in agentic experience vindicated by the scientific study of action? Versions of some of these questions can perhaps be posed from within the doxastic and telic models, but it is only the thetic model that brings them out into the clear light of day.

Secondly, the perceptual model encourages us to take a new look at perception's reach. If the sense of agency can be understood in perceptual terms how many *other* facets of human experience might also succumb to a perceptual treatment? Keeley (2002) argues that we may enjoy some form of vomeronasal perception, and Schwitzgebel and Gordon (2000) have argued that normal humans engage in echolocation. Might various forms of cognitive feelings—the experience of having a word on the 'tip-of-one's-tongue, of taking oneself to know what the answer to a question is, of taking someone to be familiar—also

qualify as perceptual? Some theorists begin their analysis of perception with the assumption that the human senses are exhausted by the famous five.¹⁴ I would suggest that this is not the best approach to take. Rather, it seems to me to be an open question—one that is partly conceptual and partly empirical—just how many senses we have. The folk notion of perception might be exhausted by the famous five, but the natural sciences employ a notion of perception that is rather more encompassing. Vision, audition, olfaction, taste and touch may be the paradigms of the senses, but there is good reason to doubt that they exhaust the category. We need a nuanced account of perception, an account that does justice not only to the famous five but also to their less celebrated siblings.¹⁵

References

- Adair, J.C., Schwartz, R.L., Na, D.L., Fennell, E., Gilmore, R.L., Heilman, K.M. 1997. Anosognosia: examining the disconnection hypothesis. *Journal of Neurology, Neurosurgery & Psychiatry*, 63: 798-800.
- Anscombe, E. 1957. *Intention*. Oxford: Blackwell.
- Bach, K. 1978. A Representational Theory of Action. *Philosophical Studies* 34: 361-379.
- Banks, G., Short, P., Martinez, J., Latchaw, R., Ratcliff, G., Boller, F. 1989. The alien hand syndrome, *Archives of Neurology*, 46: 456-59.
- Bayne, T. 2006. Phenomenology and the Feeling of Doing: Wegner on the Conscious Will, In S. Pockett, W. P. Banks and S. Gallagher (eds.) *Does Consciousness Cause Behavior?* Cambridge, MA: MIT Press, pp. 169-186.
- Bayne, T. 2008. The phenomenology of agency, *Philosophy Compass*, 3: 1-21.
- Bayne, T. & Levy, N. 2006. The Feeling of Doing: Deconstructing the Phenomenology of Agency. In N. Sebanz and W. Prinz (eds.) *Disorders of Volition*. Cambridge, MA: MIT Press, pp. 49-68.
- Bayne, T. & Pacherie, E. 2007. Narrators and comparators: the architecture of agentive self-awareness, *Synthese* 159: 475-91.

¹⁴ See for example O’Dea (this volume) and Nudds (2003 and this volume).

¹⁵ I am very grateful to Frédérique de Vignemont, Elisabeth Pacherie, Michael Schmitz, Laura Schroeter and especially Fiona Macpherson for comments on previous drafts of this chapter.

- Blakemore, S.-J. & Frith, C. D. 2003. Self-awareness and action. *Current Opinion in Neurobiology*, 13: 219-224.
- David, N., Newen, A., and Vogeley, K. 2008. The “sense of agency” and its underlying cognitive and neural mechanisms, *Consciousness and Cognition*, 17: 523-34.
- Della Sala S., Marchetti, C., Spinnler, H. 1991. Right-sided anarchic (alien) hand: A longitudinal study. *Neuropsychologia* 29(11): 1113-1127.
- Dretske, F. 1995. *Naturalizing the Mind*. Cambridge, MA: MIT Press.
- Frith, C. 1992. *The Cognitive Neuropsychology of Schizophrenia*, Hove: Lawrence Erlbaum Associates.
- Frith, C.D., Blakemore, S.J. & Wolpert, D.M. 2000a. Abnormalities in the awareness and control of action. *Philosophical Transactions of the Royal Society of London, Series B: Biological Sciences*, 355 (1404): 1771-1788.
- Frith, C. D., Blakemore, S.-J. & Wolpert, D. M. 2000b. Explaining the symptoms of schizophrenia: Abnormalities in the awareness of action. *Brain Research Reviews*, 31: 357-363.
- Gasquoine, P.G. 1993. Alien hand sign. *Journal of Clinical and Experimental Neuropsychology*, 15/5: 653-667.
- Ginet, C. 1997. Freedom, responsibility and agency, *The Journal of Ethics*, 1: 85-98.
- Gold, M., Adair, J.C., Jacobs, D.H., Heilman, K.M. 1994. Anosognosia for hemiplegia: an electrophysiologic investigation of the feed-forward hypothesis. *Neurology*, 44 (10), 1804-1808.
- Goldberg, G. & Bloom, K.K. 1990. The alien hand sign. Localization, lateralization and recovery. *American Journal of Physical Medicine and Rehabilitation*, 69: 228-38.
- Haggard, P. 2005. Conscious intention and motor cognition. *Trends in Cognitive Science*, 9(6): 290-5.
- Hallett, M. 2007. Volitional control of movement: The physiology of free will. *Clinical Neurophysiology*, 118: 1179-92.
- Heilman, K.M. 1991. Anosognosia: Possible neuropsychological mechanisms. In *Awareness of deficit after brain injury. Clinical and theoretical issues*, G.P. Prigatano & D.L. Schacter (eds.), New York: Oxford University Press.
- Hohwy, J. 2004. The experience of mental causation. *Behavior and Philosophy*, 32: 377-400.
- Holton, R. 2006. The act of choice. *Philosophers' Imprint*, 6/3: 1-15.

- Horgan, T., Tienson, J., & Graham, G. 2003. The phenomenology of first-person agency, In S. Walter and H-D Heckmann (eds) *Physicalism and Mental Causation: The Metaphysics of Mind and Action*. Exeter, UK: Imprint Academic (pp. 323-40).
- Hornsby, J. 1980. *Actions*. London: Routledge & Kegan Paul.
- Keeley, B. L. 2002. Making sense of the senses. *The Journal of Philosophy*, 99: 5-28.
- Korsgaard, C. 1996. Creating the kingdom of ends: Reciprocity and responsibility in personal relations. In *Creating the Kingdom of Ends*. Cambridge: Cambridge University Press.
- Marcel, A. 2003. The sense of agency: Awareness and ownership of action. In J. Roessler & N. Eilan (eds.) *Agency and Self-Awareness*. Oxford: Oxford University Press, pp. 48-93.
- Mellors, C.S. 1970. First rank symptoms of schizophrenia. *British Journal of Psychiatry* 117: 15-23.
- Mossel, B. 2005. Actions, control and sensations of acting, *Philosophical Studies*, 124/2: 129-180.
- Naccache, L., Dehaene, S., Cohen, L., Habert, M-O., Guichart-Gomez., E., Galanaud, D., Willer, J-C. 2003. Effortless control: executive attention and conscious feeling of mental effort are dissociable. *Neuropsychologia* 43: 1318-1328.
- Nahmias, E., Morris, S. Nadelhoffer, T. & Turner, J. 2004. The phenomenology of free will, *Journal of Consciousness Studies*, 11/7-8: 162-79.
- Nudds, M. 2003. The significance of the senses. *Proceedings of the Aristotelian Society*, 104/1: 31-51.
- O'Shaughnessy, B. 1980. *The Will (Vol. II)*. Cambridge: Cambridge University Press.
- O'Shaughnessy, B. 2003. The epistemology of physical action. In J. Roessler & N. Eilan (eds.) *Agency and Self-Awareness*. Oxford: Oxford University Press, pp. 345-57.
- Pacherie, E., Green, M. & Bayne, T. 2006. Phenomenology and delusions: Who put the 'alien' in alien control?, *Consciousness and Cognition*, 15: 566-77.
- Pacherie, E. 2007. The anarchic hand syndrome and utilization behaviour: A window onto agentic self-awareness. *Functional Neurology*, 22/4: 211-17.
- Pacherie, E. 2008. The phenomenology of action: a conceptual framework. *Cognition*, 107: 179-217.

- Peacocke, C. 2003. Awareness, ownership, and knowledge. In Roessler, J. & N. Eilan (eds.) *Agency and Self-Awareness*. Oxford: Oxford University Press, pp. 94-110.
- Peacocke, C. 2008. *Truly Understood*. Oxford: Oxford University Press.
- Roessler, J. & Eilan, N. (eds.) 2003. *Agency and Self-Awareness*. Oxford: Oxford University Press.
- Schilder P. 1953. *Medical Psychology*. New York: John Wiley & Sons, New York.
- Schwitzgebel, E. and Gordon, M.S. 2000. How well do we know our own conscious experience? The case of human echolocation. *Philosophical Topics*, 28/2: 235-46.
- Shoemaker, S. 1968. Self-reference and self-awareness. *Journal of Philosophy* 65(19): 555-567.
- Searle, J. 1983. *Intentionality*. Cambridge: Cambridge University Press.
- Searle, J. 2001. *Rationality in Action*. Cambridge, MA: MIT Press.
- Siegel, S. 2005. The phenomenology of efficacy, *Philosophical Topics*, 33/1: xxx-xx.
- Sierra, M. & Berrios, G.E. 1998. Depersonalization: Neurobiological Perspectives. *Biological Psychiatry*, 44: 898-908.
- Strawson, G. 1986. *Freedom and Belief*. Oxford: OUP.
- Tye, M. 1995. *Ten Problems of Consciousness: A Representational Theory of the Phenomenal Mind*. Cambridge, MA: MIT Press.