

# A Cognitive Account of Agentive Awareness

Myrto Mylopoulos

**Abstract:** Agentive awareness is one's awareness of oneself as presently acting. Dominant accounts in cognitive science take agentive awareness to be grounded in the states and processes underlying sensorimotor control. In this paper, I air concerns for this approach and develop an alternative. Broadly, on the approach I defend one is agentively aware in virtue of intending to act. I further argue that agentive awareness is not constituted by intentions themselves, but rather first-personal thoughts that are formed on the basis of them. I develop this proposal, highlight some of its theoretical advantages, and show how it successfully meets various challenges.

## 1 Introduction

When you engage in bodily action, you undergo various proprioceptive, kinaesthetic, and visual experiences as you guide your movements in appropriate ways. But separate from these experiences, you may also have what seems to be a distinctive awareness of yourself *as* presently acting—a sense of yourself as the author or causal source of what you do.<sup>1</sup>

Following Bayne & Pacherie (2007), let's call this *agentive awareness*.

When trying to make sense of this special kind of self-awareness, a natural question arises as to what are the psychological mechanisms or processes by which it is generated. This question is especially pressing in light of certain pathologies in which these mechanisms seem to be impaired. In some such cases, agentive awareness is absent when one might expect it to be present. Schizophrenic individuals with delusions of control, for example, appear to lack agentive awareness for some of their seemingly purposive behaviour, attributing it instead to external agents. As one patient complained, '[m]y grandfather hypnotized me and now he moves my foot up and down' (Frith et al., 2000, pp. 358; see also

---

<sup>1</sup> I will focus on bodily actions throughout this paper, leaving a treatment of agentive awareness over mental actions for another discussion (for one such treatment, see Proust, 2013, Ch.10).

Address for correspondence: 3A35 Paterson Hall, Philosophy Department, Carleton University, 1125 Colonel By Drive, Ottawa, Ontario, K1S 5B6, Canada  
E-mail: myrto.mylopoulos@carleton.ca

Spence et al., 1997; Mellors, 1970). In other cases, agentic awareness is present when one might expect it to be absent. In the condition known as anosognosia for hemiplegia (AHP), for example, patients fail to properly register their paralysis, and sometimes report illusory awareness of performing actions with paralyzed limbs (Fotopoulou et al., 2008). Berti and Pia (2006) describe the following remarkable exchange with one AHP patient who suffered from complete paralysis on the left side of her body:

Experimenter: Can you raise both your arms?

Patient [raising the right arm but not the left]: Here you are!

Experimenter: Have you raised also the left arm?

Patient: Yes.

Experimenter: Thank you. You can now put your arms down.

Experimenter: Now, could you raise your right arm again?

[P raises her right arm without any hesitation.]

Experimenter: Now could you slowly raise your left arm and

tell me when it is at the same height as the right?

Patient [after a few seconds]: Done!

Experimenter: Are you sure?

Patient: Yes! (pp. 246)

These conditions stress the need for a fuller understanding of the psychological mechanisms underlying agentic awareness and their breakdown.

A popular view is that we can make progress here by looking to low-level, sub-personal states and processes that implement sensorimotor control. A widely endorsed neurocomputational model of such control is known as the *Comparator Model* (CM).

According to the CM, whenever a motor command is computed by the motor system in the

service of some action goal, a forward model prediction of the sensory consequences of that command is generated, which is then compared with the actual sensory consequences of the movement (Wolpert, 1997; Frith, 2007). Those who seek to extend this model to account for agentic awareness suggest that when this comparison results in a sufficiently close match, a signal to the effect that the action is one's own is produced (e.g., Gallagher, 2000; Frith, 2007; Bayne, 2011b; Carruthers, 2012).<sup>2</sup> When the outcome of the comparison is a mismatch, however, a signal to the effect that the action is not one's own results instead.

Elsewhere, I have raised worries for the CM approach to agentic awareness (see Mylopoulos 2012, 2015).<sup>3</sup> As my chief aim here is constructive rather than critical, I will only briefly rehearse what I take to be its central shortcomings, and then move on to the main task at hand.

As I see it, a core problem with the CM approach to agentic awareness is that the very features of the model that make it attractive for explaining sensorimotor control simultaneously make it doubtful that it can account for agentic awareness.

This becomes evident when considering that sensorimotor control and agentic awareness can and do dissociate. This is the case in anarchic hand syndrome (AHS), a condition often caused by damage to the supplementary motor area (SMA), in which individuals perform complex movements with their contralesional limb that they are not aware of as their own actions. What is notable here, and problematic for the CM approach, is that at the level of sensorimotor control, the movements of the affected limb are well-executed—it accurately reaches for, grasps, and manipulates targets. There is no significant mismatch between the

---

<sup>2</sup> There are some who acknowledge various shortcomings of the CM, but this has largely led them to supplement or make adjustments to the account rather than abandon it entirely (e.g., Synofzik, Vosgerau, and Newen, 2008; Carruthers, 2012).

<sup>3</sup> Note that I am not disputing the strength of the CM as a model of sensorimotor control, but rather as a model for explaining agentic awareness.

forward model prediction and sensory consequences since, if there were, there would be corresponding disruptions in sensorimotor control. It is at higher levels of control— involving inhibition and guidance relative to the agent’s *intentions*—that the control deficit arises.<sup>4</sup> And yet, despite intact sensorimotor control, it is clear from their reports that AHS patients do not experience themselves as the authors of their affected limb’s movements. According to one individual, ‘of course I know that I am doing it. It just doesn’t *feel* like me’ (as reported in Marcel, 2003, 79). This flatly contradicts the main predictions of the CM approach. In short, my main concern is that the CM cannot double as a successful model of sensorimotor control *and* a successful model of agentic awareness.

There is more to be said here. But I hope that this is sufficient to motivate a search for an alternative explanation of agentic awareness that appeals to the personal-level states and processes involved in higher-levels of control—an agent’s intentions, beliefs, and perceptual states. In this paper, I will develop one such proposal by exploring the general view that agentic awareness arises in virtue of intending to act. I call this the *cognitive account of agentic awareness*.

This approach to agentic awareness has been little explored. The following will be an attempt to give it a fair shake. The plan is as follows: First, I will offer some broad support for the cognitive approach by arguing that agentic awareness arises prior to and independently of sensory feedback from bodily movement. This brings us to the natural suggestion that agentic awareness is grounded in either intentions or closely related states occurring in the pre-movement stages of action production. Adopting this suggestion, I will

---

<sup>4</sup> It is noteworthy that Blakemore et al. (2002), who did much pioneering work in developing the CM approach to agentic awareness, do not appeal to the sub-personal states of the CM to explain the experiences of AHS patients, but rather personal-level intentions (239). As the reader will see, this is broadly the direction I wish to go in as well.

argue that agentic awareness is not constituted by intentions themselves, but rather first-personal thoughts that are formed immediately on the basis of them. I will develop this account, highlight some of its theoretical advantages, and show how it can successfully meet various objections.

## **2 Defending the Independence Claim**

On the approach that I will defend, agentic awareness arises independently of sensory feedback from bodily movement; it is a kind of *default* awareness that accompanies our actions. I say ‘default’ because, on this view, it is generated automatically on the basis of one’s intention to act and it typically persists *unless* there is sensory evidence that something has gone wrong with one’s action, i.e., that one is not doing what one intends to do. This picture is at odds with models of agentic awareness on which it arises *only after* it has been confirmed, via sensory feedback, that what one is doing corresponds to one’s intention (e.g., Wegner, 2002). On the present account, sensory feedback from bodily movement may play the role of *disrupting* one’s awareness of acting by revealing that one is not doing what one intends to do, but it does not play any role in giving rise to such awareness.

Treating agentic awareness as a default accompaniment to action fits comfortably with the common observation that this type of awareness is ‘thin’ and ‘evasive’ (see, e.g., Metzinger, 2006; Tsakiris et al., 2007, 658); the kind of awareness that is more conspicuous when absent than when present (Ginet, 1990, pp. 27). On this picture, it is usually when something has gone *wrong* with our action, and agentic awareness is suddenly disturbed, that we may come to be aware that it was there in the first place. In this way, agentic awareness may be compared to the kind of awareness one sometimes has of steady background noise

in one's immediate environment. It is often only when the noise stops that one recognizes that it was previously there.

In support of this general picture, I will defend what I call the *Independence Claim (IC)*:<sup>5</sup>

(IC): Agentive awareness arises prior to and independently of sensory feedback from bodily movement.

If the IC is true, then this helps to establish the default character of agentive awareness in that confirmation on the basis of sensory feedback that one is doing what one intends to do is not needed for it to arise. But we should not expect that evidence for the IC is easy to come by. After all, sensory feedback from bodily movement typically follows almost immediately after action initiation. If I initiate the raising of my arm, almost straightaway I sense my muscle contractions and changes in the positioning of my limbs. So identifying cases in which agentive awareness clearly arises *prior* to such feedback is a challenge. Still, there are some cases that together lend support for the IC. I will discuss two: (i) deafferentation, and (ii) subjective timing studies.

Deafferentation is a condition that results in significant impairments to touch and proprioception. There are cases of deafferented individuals who have lost almost all sense of proprioception, the most well-known of which are Ian Waterman (see Cole, 1995) and GL (see Fournier et al., 2002). At the age of 19, Waterman was struck with a virus that led to deafferentation from the neck down. At first, he lost all control of his body, being unable to sit or stand up, walk, feed himself, or interact with ordinary objects. But after years of rigorous practice and rehabilitation, he is now able to guide his body remarkably well with

---

<sup>5</sup> Note that what I am disputing in this section is the role of sensory feedback in generating agentive awareness, rather than the more general claim that agentive awareness is a form of sensory awareness, as some have argued (e.g., Bayne, 2011b). For arguments against that view, see Mylopoulos (2015).

the help of vision (Cole, 1995). GL, at the age of 31, suffered carbon monoxide poisoning resulting in sensory neuropathy. She was left with deafferentation from her nose to her toes (Guillaud, Simoneau, & Blouin, 2011). In her case, too, intensive training and rehabilitation has helped her regain visually-guided control of her body.

Are deafferented individuals like Ian Waterman and GL agentively aware when they perform actions in the absence of sensory feedback? There is some evidence that they are. For example, commenting on GL's performance in an experimental task, de Vignemont and Fournier (2004) note that: 'She did not pretend to have moved while she did not, like anosognosic patients, nor did she make the same movement twice believing that she had not moved yet' (pp. 13).

Still, there is room for skepticism here. For instance, Carruthers (2012) writes:

... it seems these patients do not experience a sense of agency over actions they do not see, at least for the case of conversational gestures. Both of these patients [IW and GL] gesture when talking. GL was shown a video of her doing so. On seeing this she reported that she did not feel that she controlled the movements at the time they occurred... (pp. 33)

Furthering this line, Carruthers describes a study in which Waterman was asked to narrate the plot of a cartoon while wearing a blindfold (Cole, Gallagher, & McNeill, 2002). Twenty seconds in, after he had already produced fourteen conversational gestures, Waterman reportedly said, '... and I'm starting to use my hands now...', suggesting that he had not been aware of doing so previously.

I'm not convinced that the evidence in question undermines an appeal to cases of deafferentation in support of the IC. First, GL's report of not feeling in control of her movements at the time that she made them is compatible with her not being aware of making them at all, something that is not uncommon when we engage in highly automatic or habitual movements, like conversational gestures, that involve minimal guidance from conscious intentions. Although in the early stages of their deafferentation, the conversational gestures of Waterman and GL were forced, effortful, and deliberate, after much training they have each succeeded in making them routine and effortless (Cole, Gallagher, & McNeill, 2002). In addition, Waterman's report that he is starting to use his hands suggests that at that point he *does* have agentic awareness, despite the absence of proprioceptive and visual feedback. Carruthers notes this as well, and suggests that '[m]ore likely this report comes from knowledge of his intention to begin moving. In other words this report is a result of I.W.'s theory of mind and not his sense of agency' (pp. 34). But the very question I am interested in is whether agentic awareness might indeed be based on something akin to knowledge of one's intention to begin moving, so in the absence of an independent reason to think this is not what is happening here, I take cases of deafferentation to serve as a reasonable source of support for the IC.

Combined with cases of deafferentation, there is experimental evidence in favour of the IC that stems from subjective timing studies. Here we can appeal to the classic work led by the neuroscientist Benjamin Libet—and the subsequent replications and refinement of that work (e.g., Haggard & Eimer, 1999). Libet et al. (1983) set out to determine the temporal relationship between one's willing and initiating an action, and one's awareness thereof. They asked six participants each to perform a series of spontaneous simple actions. More specifically, they were instructed in each trial to perform a 'quick, abrupt flexion of the



fingers/and or wrist' (625) with their right hand, at a time of their own choosing.

Participants were seated facing a specialized clock, around which a dot would revolve. They were hooked up to an electroencephalogram (EEG), which measured their brain activity, as well as an EMG, which measured their muscle activity. They were told to report the time at which they first became aware of deciding or having the 'urge' to move based on the position of the dot on the clock (the so-called 'W-judgment'). They were also asked to report the time at which they first became aware that they had *actually* moved (the so-called 'M-judgment').

Most of the discussion (and controversy) surrounding this study, and others like it, has focused on the striking finding that, though participants report being aware of the decision to move approximately 200 ms, on average, before the onset of movement, a neural event known as the 'Readiness Potential (RP)' that many take to be the neural signature of action initiation, begins to take shape around 350 ms *prior* to this (see Schurger, Sitt, & Dehaene, 2012; Schurger, Mylopoulos, and Rosenthal, 2016; Mele, 2009 for discussion of this result). While interesting, I wish to set this result aside for now, and focus instead on the M-judgment, which reflects the time at which participants *first* come to have an awareness of acting.

The findings here are revealing: participants report being aware of acting an average of 86 ms *prior* to the onset of muscle activation as measured by the EMG. Moreover, this result has been widely replicated; anticipatory awareness of action has since been reported in a number of subjective timing studies using similar paradigms (e.g., Haggard & Eimer, 1999; Haggard, Newman, & Magno, 1999; Lau et al., 2004). This strongly suggests that agentic awareness arises independently of and prior to any sensory feedback from bodily movement.

Still, this robust finding is only good evidence in favour of the IC if the M-judgement reports are reliable. And one might be doubtful that they are. One reason to worry might be based on the results of an elegant study by Lau, Rogers, & Passingham (2007), which suggest that participants' M-judgments are sensitive to events occurring *after* the movement is performed. Using a Libet-style subjective timing paradigm, Lau et al. asked participants to perform spontaneous button presses, and then to indicate the time at which they became aware of acting. In half the trials, transcranial magnetic stimulation (TMS) was administered over the pre-SMA, and in the other half, a 'sham TMS' was activated. The purpose of the sham TMS trials was to prevent participants from knowing during which trials the pre-SMA area of their brain was actually being stimulated. The real or sham stimulations were applied either simultaneously with the participant's action or 200 ms afterwards.

Lau et al. found that TMS stimulation over the pre-SMA induced a forward shift in the reported awareness of the onset of action even when it was administered 200 ms *after* the action had begun. On these trials, there was a 9 ms delay, on average, in the reported timing of the action onset relative to that reported in the sham TMS trials. This effect, though small, was found to be significant: the application of TMS after movement onset had the result of distorting participants' M-judgments (for similar TMS manipulations to M-judgment, see also Haggard, Clark, and Kalogeras, 2002; Haggard and Magno, 1999).

It may be tempting to interpret these results as casting doubt on the reliability of M-judgments. Perhaps, one might argue, the results suggest that these judgments do not accurately reflect the time at which participants first become aware of acting, but that they are retroactive judgments that are based on cues occurring only *after* movement onset.

There are reasons to resist this temptation. First, it does not follow from the fact that participants' M-judgments can be interfered with via TMS administered after action onset

that, in the absence of such interference, they are not reliable. Lau et al. (2007) point out that the TMS is likely to interfere with signals corresponding to action preparation, which could mean that since the typical cues on which agentic awareness is based are degraded, participants rely more on sensory feedback from bodily movement than they otherwise would, thus resulting in the judgement that they moved later than they did. Finally, it is worth stressing that even with the TMS-induced delay, the M-judgement is *still* prior (-41 ms on average) to the onset of bodily movement, so in this way the IC is not threatened.

The two foregoing lines of evidence point convincingly, I think, to the involvement of states occurring prior to and independently of sensory feedback from bodily movement in generating agentic awareness. But there is still a question as to *which* pre-movement states are implicated. Some (e.g, Prinz, 2007) have argued that the relevant states are sensory states, in particular forward model predictions, as these occur prior to action and represent the sensory consequences of movement. But forward model predictions are rapidly computed, short-lived states involved in fine-grained control of bodily movement, and there is a question as to whether we are ever conscious of them. Perhaps even more worrisome, we have no reason to think that being aware of the sensory aspects of bodily movement, as specified by the forward model prediction, should give rise to *agentic* awareness as opposed to awareness of *passively* moving, which would also involve awareness of these sensory aspects. Why should the fact that in one case, one is predictively aware of the sensory aspects of a movement, while in the other, one is aware of them as they occur, determine whether one is agentively versus passively aware of the movement in question? I suggest that a more reasonable conclusion to draw from the IC is that *we are agentively aware in virtue of the very intentions that initiate and guide our actions*. This is the cognitive approach to agentic awareness.

There are two ways to understand the ‘in virtue of’ relation that I am appealing to here. One might treat it as a constitutive relation, as when there is water in the glass in virtue of there being H<sub>2</sub>O in the glass. Or one might view it as a causal relation, as when there is smoke in the kitchen in virtue of there being a fire in the kitchen. Which sense is operative here? A constitutive reading would have it that intentions themselves are the states that ground agentic awareness. What it is to be agentially aware is to intend to act. I call this the *volitional account*. A causal reading, by contrast, would have it that agentic awareness is somehow causally related to an intention to act, though not constituted by it. I will argue in the last section that this is the correct view. First, though, I explain my reasons for rejecting the volitional account.

### **3 The Volitional Account (And Why It Doesn’t Work)**

It is common to distinguish between two types of intention that differ in terms of their temporal and functional dynamics (e.g., Searle, 1983; Brand, 1984; Bratman, 1987; Mele, 1992; Pacherie, 2008). *Distal intentions* concern future actions and are key players in planning, practical reasoning, and interpersonal coordination (see Bratman, 1987). What I call *executive intentions*, on the other hand, are concerned not with future action, but with what to do *now*. Ginet (1990) nicely captures the character of such intentions when he remarks that, ‘they do not plan ahead, not even very slightly. They do not *plan* at all; they *execute*’ (33). Executive intentions, then, are the psychological states that initiate and guide present action; they are the most proximate mental causes of action.

Distal intentions are formed too far in advance of action onset, so they are plainly not suitable states for making one aware of oneself as presently acting. We must instead look to executive intentions and see how they fare. From here on, when I speak of intentions, it is the executive variety that I have in mind. Could such states constitute agentic awareness?

A classic proponent of this view is John Searle (1979, 1983) (for similar proposals, see also Ginet, 1990; Marcel, 2003; Peacocke, 2007; Mandik, 2010), who identifies what he calls *experiences of acting* with his own brand of executive intention, which he calls *intention in action*. Experiences of acting, according to Searle, are just intentions in action that possess characteristic phenomenal properties: ‘... the experience of acting just is the intention in action. The only difference between them is that the experience may have certain phenomenal properties that are not essential to the intention’ (265). Since phenomenal properties, for Searle and others, are possessed by conscious states only, we are left with the view that one’s awareness of acting is constituted by one’s conscious intention to act.

This picture is simple and intuitive. After all, given the role of executive intentions in initiating and guiding action, it is natural to suppose that these states provide one with information that one is presently engaged in action. There are at least two major problems here, however. The first is exposed once we consider the content of one’s awareness of one’s mental states. Many types of mental state exhibit both representational content (e.g., *there is milk in the fridge*) and a mental attitude directed towards that content (e.g., *believing, doubting, or hoping* that there is milk in the fridge). When such mental states are conscious, one is aware not only of their representational content, but also their mental attitude. I cannot be aware of my belief, doubt, or hope that there is milk in the fridge without being aware of it *as* a belief, doubt, or hope.

Now, this may seem rather obvious, but it spells trouble for the constitutive account. To illustrate the problem, let us consider Searle’s (1979) claim that when I have an intention in action to raise my arm, ‘[t]he experience of acting is *of* the movement of my arm...’ (268). This corresponds to the representational content of the intention in action, which Searle takes to be *that my arm goes up by way of this intention in action*. Compare this to the case of

believing that there is milk in the fridge. Suppose that having a conscious belief gives rise to an experience in the same way that Searle takes conscious intention in action to give rise to an experience. What is this experience an experience of? It is not an experience of there being milk in the fridge, but an experience of *believing* this to be the case. Likewise, we should not expect that a conscious intention in action results in an experience or awareness of moving one's arm, but in the experience or awareness of *intending* to move one's arm.<sup>6</sup> Put more generally, a conscious executive intention to  $\phi$  gives rise to an awareness of *intending* to  $\phi$ , rather than an awareness of  $\phi$ -ing. But if so, then the constitutive view does not explain awareness of *acting*, it explains awareness of *intending* to act.

Recently, Kriegel (2015) has offered what might be thought of as a way out of this worry. He argues by way of analogy with visual experiences that an experience of acting just *is* an experience of successfully trying to act. Though he focuses on trying, let's consider whether his argument can be generalized to apply to successful executive intentions as well. He writes:

We certainly experience ourselves as acting, or in other words as successfully trying to do something. But we also experience ourselves as seeing the world, that is, as in a good case of visual experience. We do not normally experience ourselves as hallucinating or as being in a state that might be either a seeing or a hallucinating. All the same, our experience is in fact a state which might be either a seeing or a hallucinating.

When it is a seeing, the phenomenology is veridical, and when

---

<sup>6</sup> Note that the same problem arises if we treat the content of the intention in action as *I am  $\phi$ -ing*, in the way that McDowell (2010) suggests that we do.

it is a hallucinating it is nonveridical. Likewise with trying: when it is successful, our experience of ourselves as acting is veridical, and when it is unsuccessful, nonveridical. It remains that nothing in the conative experience itself guarantees its success, just as nothing in a visual experience guarantees its veridicality. So the experience itself is just a trying (90).<sup>7</sup>

Kriegel (2015) is right to point out that there is nothing in the experience of intending that indicates the success or failure of the intending. But notice that in the visual case, the hallucinatory awareness and the veridical awareness are still both such that one is aware of *seeing* something. So, if the analogy is strictly to hold, *both* successful intendings and unsuccessful intendings should yield a corresponding awareness of *intending* to do something. This is precisely the problem, though. We do not want to explain awareness of intending, but awareness of doing.

One might reply here that what is important is that the two types of awareness—of unsuccessfully intending to  $\phi$  and of successfully intending to  $\phi$ —are subjectively indistinguishable from an awareness of  $\phi$ -ing. But the Libet-style subjective timing studies discussed earlier suggest that one's awareness of intending to act *is* subjectively distinguishable from one's awareness of acting. Participants report their awareness of their intentions and actions as occurring at separate times, indicating that they can discriminate between them. There is a significant temporal gap between the awareness of intending to act as expressed in the W-judgment and the awareness of engaging in action, as expressed in the

---

<sup>7</sup> Indeed, as Bayne (2008) points out, Searle (1983) seems to endorse this very position in his discussion of William James' anaesthetized patient (89).

M-judgment.<sup>8</sup> Participants do not take the intention to start acting to be subjectively indistinguishable from their action.

But perhaps it is wrong to view the states tracked by the W-judgement as intentions. Perhaps, given their spontaneous and unplanned nature, it is better to view them as *urges*, so that it is awareness of *urges* that is subjectively distinguishable from awareness of acting, but not awareness of intentions. But some intentions are spontaneously acquired as well—think of dodging an unexpected object or Searle’s famous example of jumping out of your chair as you ruminate on a philosophical puzzle. In addition, the awareness in question is one of having *settled* or *decided* on acting at that moment, which is characteristic of executive intention, and not merely an awareness of being pulled or tempted towards action, which is characteristic of urges. So I think there is good reason to view the W-judgement as a report of an intention to move, and not just a strong urge to do so.<sup>9</sup>

In support of the subjective distinguishability between intending and acting, we can also consider the everyday case of engaging in an action that falls short of satisfying one’s intention. Suppose you are crouched down, attempting to lift a heavy box, but so far haven’t mustered the strength to make it budge. You are here aware of intending to lift the heavy box, but not yet aware of lifting it. Awareness of intending to  $\phi$  and awareness of  $\phi$ -ing are importantly different, and it seems that the constitutive view can only help to explain the former.

---

<sup>8</sup> As noted earlier, there is some controversy surrounding the correct interpretation and significance of the W-judgment, but no one to my knowledge disputes that the awareness of the intention to move occurs prior to the awareness of acting as indicated by the participants’ W- and M-judgments.

<sup>9</sup> In a recent paper, Shepherd (2016) convincingly argues that experiences of trying, which he views as experiences of directing effort towards the satisfaction of an intention, are separate from the experience of acting. One possibility, then, is that when properly attended to, the preparatory stages of action involve at least two distinguishable experiences: one of intending and one of trying to act, followed by an experience of acting.



Before moving on, it's worth highlighting a second difficulty with the constitutive view. To illustrate, I must say something about the notion of direction of fit (Anscombe, 1966; Searle, 1983). A state's direction of fit can be profitably understood as being determined by the functional role that it plays within one's mental economy (cf. Bayne 2011a). We can distinguish between *thetic* (also sometimes known as *mind-to-world*) and *telic* (also sometimes known as *world-to-mind*) directions of fit (Humberstone, 1992; Bayne, 2011a; Pacherie, 2015). Mental states with a *telic* direction of fit are those with the function of driving changes in the world. They tend to cause the world to be how they represent it as being. As Searle observes, we evaluate the success of *telic* states by determining whether they have been fulfilled or realized, that is, whether they have succeeded, rather than failed, in causing the world to match up with their content. Intentions and desires are often thought of as paradigmatic *telic* states. By contrast, the functional role of *thetic* states is not to drive changes in the world, but to represent the way the world is. When assessing their success, we do so by evaluating them for truth; we determine whether they accurately represent the way the world is. Beliefs and perceptions are paradigmatic *thetic* states.

It seems that the proponent of the constitutive view must treat experiences of acting as having *telic* structure, since they are constituted by intentions. In support of this construal, Searle (1979) writes: '... in the case of the experience of acting... [i]f I have this experience but the event doesn't occur we say such things as that I *failed* to raise my arm, or that I *tried* to raise my arm but did not succeed' (262), thus pointing out that we tend to use the evaluative terms linked to *telic* states when assessing agentive experiences.

Searle is correct that in such cases one might indeed say of *oneself* that one failed to raise one's arm, or that one tried to raise one's arm, but did not succeed. And one might likewise say that one's *intention* or *trying* to raise one's arm was unsuccessful. But one would

plainly *not* say that the *experience* or the *awareness* of raising one's arm failed or was unsuccessful. Rather, the appropriate evaluation here concerns whether or not the experience or awareness is *accurate* or *veridical*. And this, in turn, tells us that the correct way to view such awareness is as being grounded in thetic, not telic states, and so not intentions.<sup>10</sup>

But perhaps this is too quick. Recently, it has been suggested by Bayne (2011a) that agentic awareness is a complex, dynamic experience unfolding in time (for similar views, see also Pacherie, 2015; Shepherd, 2016), that may have *both* telic and thetic structure, and thus is best construed as being grounded in 'pushmi-pullyu representations' that both direct and describe an agent's action, just as a bee dance both directs other bees to the source of the nectar and describes the nectar as being present at a certain location (see Millikan 1996). Bayne (2011a) elaborates on this proposal as follows:

According to this dynamic conception of agentic experience, we might think of agentic experiences as having an overall pushmi-pullyu structure, with the telic components predominating early in the representational process and thetic components coming to the fore later in the representational process. The overall experiential state, however, can be assigned both realization and veridicality conditions (229).

---

<sup>10</sup> One may worry here that when we apply such evaluative terms in these contexts, we are really referring to beliefs accompanying agentic awareness, and not the awareness itself. But if so, then we should expect this to apply to other types of awareness more broadly, and it does not. When one's visual awareness presents the two lines in the Müller-Lyer illusion as being of unequal length, despite one's belief that the lines are of equal length, clearly it is one's visual awareness that is non-veridical and inaccurate, not one's accompanying belief.

To be sure, action execution involves a complex set of interacting representations (e.g., intentions, beliefs, proprioceptive, and visual states), some of which are telic (especially those occurring early on in the preparatory stages) and some of which are thetic (especially those occurring later on as the action unfolds). And it may be that these states, when conscious, give rise to corresponding experiences that capture their telic or thetic character. But I am skeptical of the *further* claim that action involves an *overarching experience* or *awareness* of acting that itself possesses a dual direction of fit. First, experiences are individuated by appeal to their content. So we must keep distinct the experiences corresponding to the various telic and thetic states involved in action execution, which differ in their content. There is not one corresponding experience that collectively subsumes these states, but a cluster of experiences succeeding each other in time. Second, at the core of agentic awareness is the self-attribution of an action in which one is presently engaged. That is what I have set out to explain in this paper. And this awareness does *not* play a role in driving the very action that is self-attributed, it is a descriptive awareness accompanying that action. Thus, we have good reason to view such awareness as being thetic in structure, and so not constituted by a conscious intention. There may very well be other experiences associated with an action as it unfolds, reflecting both telic and thetic intentional structure, but these do not, in my view, constitute agentic awareness.

#### **4 Agentic Thoughts and Agentic Awareness**

In line with the foregoing considerations from sections 2 and 3, I propose that the states that constitute agentic awareness are not intentions themselves, but states that are formed *on the basis* of them. In particular, I propose that when one forms an executive intention to perform some action, one regularly has a corresponding thought that one is performing that very action, and it is by way of such thoughts that we are agentially aware. What it is to be

aware of oneself as  $\phi$ -ing is, quite simply, to have the thought that one is  $\phi$ -ing. These thoughts are personal-level states, by which I mean that they are attributable to the agent as a whole, rather than a subsystem of the agent, they are accessible to consciousness—though as we will see below, we are not *typically* aware of them—and they are inferentially integrated within an agent’s mental economy.

The robust link between executive intentions and agentive thoughts is established on the basis of the high reliability with which executive intentions cause the actions that they represent. Of course, we sometimes blunder, misstep, or otherwise fail to pull off what we executively intend to do, but the vast majority of the time, we are successful. Executive intentions are formed in the moment of action, and tend to specify actions in terms of the bodily movements required to carry them out, which are typically robustly well-learned sequences. There is thus a strong causal connection between executive intentions and the actions that they specify. Sensitivity to this connection over time results in what might be viewed as a mental habit, whereby an agent automatically forms an agentive thought to the effect that she is doing what she executively intends to do. Once this mental habit is acquired, sensory feedback from the bodily movement itself is not required for the formation of such thoughts, though as mentioned, it may lead one to abandon an agentive thought if one recognizes that one is not actually doing what one intends to do.<sup>11</sup>

---

<sup>11</sup> It is worth noting here certain similarities between the present proposal and the cognitive or doxastic accounts of *self-knowledge* of action. I take such accounts to have as their primary explanatory goal how best to characterize the epistemic access that we have to our intentional actions. In other words, how it is that we know what we are intentionally doing in the distinctive way that we do (see Schwenkler 2012 for a thorough treatment of the main positions in this debate). For instance, on Falvey’s (2000) view, one has ‘knowledge in intention’ of what one is doing which is given by a judgment that expresses one’s intention. The chief differences between such accounts and my own are the following: (i) I am not concerned with the justification or warrant ascribable to agentive thoughts, but rather the psychological mechanism that gives rise to them, and (ii) I am interested in explaining the

An immediate worry for the present proposal is that it seems to be vulnerable to one of the main objections I raised against the volitional view, and is therefore a nonstarter. In particular, one might be concerned that a conscious agentic thought will yield awareness of *thinking* that one is acting, not awareness of acting. But this is not what we set out to explain. So once again it seems that we are failing to give an account of the target phenomenon.

This concern relies on the erroneous assumption that agentic thoughts are always conscious. But they are rarely so. When we act, we are aware of aspects of our bodily movements and perhaps some of our intentions, but we are not typically aware of our agentic thoughts. Agentic thoughts make us aware of ourselves as acting even if they are not conscious themselves. Indeed, this allows the view to accommodate a distinction that is commonly drawn between a ‘reflective’ (or ‘detached’) sense of agency and an ‘immersed’ or ‘minimal’ sense of agency (see, e.g., Pacherie 2008, p.195; Marcel 2003; Gallagher 2007).<sup>12</sup> We might experience reflective agentic awareness when we are performing slow, deliberate actions that require conscious focus on what we are doing, e.g., pouring liquid into a cup from a heavy bottle. In these cases, we are explicitly aware of what we are doing: our agentic thoughts describing our actions are conscious. By contrast, *most* of our everyday actions—our more mundane, habitual ones—like brushing our teeth and tying our shoelaces, may merely involve an implicit awareness of what we are doing supplied by agentic thoughts that we have no accompanying awareness of.

---

subjective character of agentic awareness. This is not something that theories of agentic self-knowledge care to account for.

<sup>12</sup> Here I depart from the popular view that the minimal sense of agency involves non-conceptual states, since agentic thoughts are plainly conceptual states. But one main motivation for this view is that it would provide a basis for distinguishing between a minimal vs. reflective sense of agency. I have offered another way for doing so here.

Another concern for this proposal is that it cannot accommodate the *phenomenology* of agency. Many are attracted to the view that there is such a distinctive phenomenology.<sup>13</sup> As Horgan, Tienson, and Graham (2003) put it, ‘... there is “something it is like” to behave in a way that constitutes voluntary action, something phenomenologically distinctive that incorporates but goes beyond the phenomenology of one’s own bodily motion’ (323).

If we accept that there is a distinctive phenomenology of agency, we might worry that an account on which agentic thoughts are the states that provide us with agentic awareness is doomed to fail. After all, phenomenal character is often assumed to be a feature of sensory or affective states only, e.g., smells, pains, and fears. Thoughts, by contrast, are frequently taken to be phenomenally ‘nude’, i.e., to exhibit no proprietary phenomenal character of their own. On the orthodox view, whatever phenomenology is associated with thoughts is the result of *accompanying* states, such as auditory states of ‘inner speech’ or visual imagery, but not thoughts themselves.

More recently, however, there has been growing support for the view that conscious thoughts do, in fact, have their own proprietary *cognitive* phenomenology (see, e.g., Pitt 2004; for a rich overview of the present debate concerning cognitive phenomenology, see the essays in Bayne and Montague 2012). But even if we adopt this position, it does not help us here. For the issue with which we are concerned is not whether there is something it is like to have a conscious thought, which is the claim that proponents of cognitive phenomenology are eager to defend, but whether having an agentic thought is sufficient for

---

<sup>13</sup> It is worth noting that the claim that there is a positive, distinctive phenomenology of agency is typically taken for granted and rarely argued for. An alternative view would be that there is no such thing, but only a phenomenology of alienation or disruption of our actions when something goes wrong (see Prinz 2012, pp. 237 - 239). In this paper, I am working with the assumption that there is a positive phenomenology of agency, but I flag that there is an important question here worth addressing on another occasion.

there to be something it is like *for one to act*. Indeed, the cognitive phenomenology debate is about conscious thoughts exclusively. But as mentioned, agentive thoughts need not be conscious in order to make one agentively aware. They may do so nonconsciously. So the question is whether, even as nonconscious thoughts, they give rise to agentive phenomenology.

One way for the cognitive account to accommodate agentive phenomenology makes direct contact with consciousness research. On higher-order thought (HOT) theories of consciousness, what it is to be in a conscious mental state is to represent oneself as being in that state by way of a higher-order thought (HOT) (see Rosenthal 1985, 2005; Gennaro, 1996; 2012). Proponents of this view, which has enjoyed increasing empirical support recently (see Lau & Rosenthal 2011; Lau & Brown forthcoming), tend to hold that being suitably aware of oneself as being in a mental state is sufficient for there to be ‘something it is like’ for one to be in that state. The idea is that how things subjectively seem to us is a matter of what mental states we represent ourselves as being in via HOTs. Here is Rosenthal (2011) on this point:

As many, myself included, use that phrase, there being something it’s like for one to be in a state is simply its seeming subjectively that one is in that state. [...] And on that construal of ‘what it’s like’, the theory does hold that a [higher-order thought (HOT)] is sufficient for there to be something it’s like for one to be in the state the HOT describes... (pp. 433-434)

The proponent of the cognitive account might hold something analogous to be true in the case of agentive thoughts. On this proposal, just as having a HOT that represents

oneself as being in a mental state is sufficient for there to be something it is like to be in that mental state, so too having an agentive thought that represents oneself as acting is sufficient for there to be something it is like to act. And just as it is not necessary for HOTs themselves to be conscious in order to give rise to the subjective experience of seeing red, or feeling pain, it is not necessary for agentive thoughts to be conscious in order to give rise to a subjective experience of acting. This provides a promising way of making sense of how agentive thoughts can accommodate the phenomenology of agency.

Indeed, the cognitive account may be in a unique position to explain what some take to be a central feature of the phenomenology of agency, namely the experience of ‘self as source’. Here is Horgan (2012) describing this core element:

Suppose that you deliberately do something—say, holding up your right arm with palm forward and fingers together and extended vertically. What is your experience like? To begin with, there is of course the purely bodily-motion aspect of the phenomenology—the what-it’s-like of being visually and kinesthetically presented with one’s own right hand rising with palm forward and fingers together and pointing upward. But there is more to it than that, because you are experiencing this bodily motion not as something that is ‘just happening,’ so to speak, but rather *as your own action*. You experience your arm, hand, and fingers as being moved *by you yourself*; this is the what-it’s-like of *self as source* (64).

Agentive thoughts are first-personal thoughts to the effect that ‘I am  $\varphi$ -ing’, where ‘ $\varphi$ ’ is substituted with an appropriate action description “inherited” from the content of one’s



executive intention. The deployment of the essential indexical in their content makes it so that agentive thoughts explicitly represent *oneself*, as such, as the agent of the action in question. More specifically, if I have the thought that I am  $\varphi$ -ing, the essential indexical makes it the case that my thought represents as  $\varphi$ -ing the thinker of that very thought. In this way, an intimate link is forged between the subject of agentive awareness and the agent of the action, which might very well yield a robust sense of oneself as the source of that action.

This view can also help explain why it is that our awareness of our own actions seems to be direct and immediate. When you act, you do not come to be aware of what you are doing by consciously observing your body or consciously inferring what action you are engaged in. On the present view, this is explained by the fact that agentive thoughts are not formed on the basis of any conscious inference or observation—though they start out being formed in this way before the mental habit is acquired—but are rather associatively linked with executive intentions. Moreover, executive intentions themselves often fail to be conscious. Subjective timing studies may give the opposite impression, but they involve a rather unusual demand. When we perform actions under ordinary circumstances, we are not typically aware of or attending to the executive intentions that trigger them. Since the states on which they are based are often not present in conscious experience, this lends a further degree of subjective immediacy to agentive thoughts and the awareness of acting with which they provide us.

Importantly, the present account may also fit well with our understanding of what is happening in certain pathological conditions relating to agentive awareness. I focus here on anosognosia for hemiplegia (AHP), the disorder characterized by denial of paralysis on the patient's contralesional side. (For discussion of how the present view might help shed light on delusions of control in schizophrenia and anarchic hand syndrome, see Mylopoulos

2015). There is good evidence that despite being unable to move the paralyzed parts of their bodies, individuals with AHP can nonetheless form executive intentions to do so as well as issue the relevant motor commands (see Fotopoulou et al., 2008; Jenkinson & Fotopoulou, 2010). Among the strongest pieces of evidence for this is that they exhibit so-called *bimanual coupling effects*. In a healthy case, when one is asked to draw a straight line with one hand and a circle with the other, there is an interference effect so that one ends up drawing an oval with the hand that is supposed to be drawing a line. Garbarini et al. (2012) found that when AHP patients are asked to perform the same task, namely to draw straight lines with their unaffected hand and a circle with their affected, i.e., paralyzed hand, the unaffected hand draws an oval despite no movement from the affected limb. Such interference effects were not found for hemiplegic patients without anosognosia. Findings like this strongly suggest that when AHP individuals are asked to perform an action with their paralyzed limb, they form the executive intention to do so, perhaps because they do not believe that they cannot, and that this action representation is what is driving the interference effect.

If AHP patients form executive intentions to move their paralyzed limbs, then they may continue to automatically form agentive thoughts on the basis of these intentions. In regular cases in which one recognizes that one is not doing what one forms an executive intention to do, one will abandon or update the corresponding agentive thought accordingly. But individuals with AHP, not being properly sensitive to their paralysis, may fail to do so, and their agentive thoughts may thereby persevere.

This is consistent with an interpretation of AHP that some in the clinical literature currently endorse (e.g., Vuilleumier, 2004). Indeed, a recent study by Vocat, Saj, and Vuilleumier (2013) titled, ‘The riddle of anosognosia: Does unawareness of hemiplegia involve a failure to update beliefs?’ is among the first to offer some direct empirical support

of this hypothesis. Vocat et al. (2013) had nine AHP patients as well as eleven healthy controls perform a task in which they were required to guess ten target words (e.g., ‘cow’). Each target word was accompanied by five verbal clues (‘I am sometimes black and white,’ ‘I am an animal of the female gender’), which gave hints as to the semantic or syntactic features of the word. These clues were presented one after the other, with each successive clue providing more information than the previous one. While the first clue left many possible answers, the last clue was meant to leave no doubt as to the identity of the target word. After each clue, the participant would make a guess as to what the target word was, and rate their level of confidence in their guess. They were given no feedback as to the accuracy of their guess. The identity of the target was revealed after each participant’s fifth and final guess and confidence rating.

The results were that AHP patients, compared with controls, showed abnormally high confidence in their guesses after the first three clues. Indeed, their confidence ratings after the *first* guess ( $M = 5.7$  on a scale from 0 – 8) were comparable with the ratings that controls gave after their *fourth* guess ( $M = 5.6$ ). Perhaps even more striking, AHP patients failed to revise their previous guesses even when new information was presented that conflicted with them. The experimenters note that, ‘[t]hey typically preferred to find “non-obvious” but “plausible” connections between the false word provided on a preceding trial rather than reject their current beliefs and make a new guess’ (1777). This study, therefore, offers some support for the view that AHP patients have difficulty updating their beliefs in light of new evidence. Applied to agential awareness, this result might be taken as support for the view that such individuals fail to revise or update their agential thoughts in light of new evidence that they are not acting—and because of this they continue to have agential awareness of performing actions that they do not actually perform.

## 5 Objections and Replies

In this closing section, I respond to a number of challenges that one may be tempted to raise against the view that I have proposed.

### 5.1 Agentive Thoughts Are Not Necessary for Agentive Awareness

One difficulty that this account might seem to face is that agentive thoughts do not appear to be necessary for agentive awareness. Here is Peacocke (2007) raising a version of this challenge:

It may seem to the unfortunate person whose arm is, unbeknownst to him, severed in a car accident that he is moving his arm, even though he has no sensation in it. This seeming has a false content. The seeming, just like a visual illusion, can persist after the subject knows his unhappy situation. In my view, action-awareness should not be identified with any kind of belief, whether first- or second-order (359).

The worry here is that one can have agentive awareness despite believing that one is not performing an action. Indeed, this seems to take place in certain so-called ‘phantom limb’ cases, in which people continue to experience the presence of their amputated limbs and sometimes even feel that they can move them at will (see, e.g., Ramachandran & Hirstein, 1998). And yet these individuals are not delusional; they know that they are not actually moving their missing limbs. But if so, then how could their illusory agentive awareness still be the result of agentive thoughts to the effect that they are acting? Wouldn’t such thoughts conflict with what they know?

This scenario only poses a problem if we accept that one cannot simultaneously be in two occurrent cognitive states with contradictory contents. In other words, we must accept that one cannot simultaneously have an agentive thought that one is moving one's limb at will, which gives rise to an awareness of acting, while also believing that one is *not* moving one's limb at will. But there are familiar cases we can point to in which cognitive states co-exist despite clashing with each other in this way. First, consider a case in which you have a pre-reflective intuition that P, but your reflective theoretical commitments lead you to believe that not-P. For example, perhaps, intuitively, you hold that a certain action is immoral—it *seems* like the wrong thing to do—but on the basis of careful ethical reasoning you have formed the belief that despite its seeming intuitively wrong, the action in question is not actually so. Despite your belief that not-P, your intuition—a cognitive state—may persist in making it *seem* to you that P is the case. Second, consider feelings of déjà vu, which are sometimes classified as metacognitive states. I may enter a room and have the strong sense that I have previously encountered the scene before me. And this sense might persist despite my knowledge that this isn't so. These are both cases where cognitive states are encapsulated to a degree from other cognitive states with conflicting content. This may be precisely what is happening with persistent agentive thoughts, which are automatically generated on the basis of executive intention, and conflicting beliefs, which result from a different source, as in 'phantom limb' cases.

## **5.2 Agentive Thoughts Are Not Sufficient for Agentive Awareness**

One may similarly worry that agentive thoughts are not sufficient for agentive awareness. Bayne (2011b) supplies the following case involving AHS to illustrate this concern:

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 [sic] 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 agentic experience? That seems highly implausible (360).

In response to this worry, I stress that agentic thoughts are not just thoughts of any type, they are thoughts with a particular aetiology: they are products of a mental habit, and they arise on the basis of executive intentions. As mentioned, they do so in a way that does not seem to rely on any observation or inference, as do the thoughts to which Bayne (2011b) is appealing in the hypothetical scenario. They provide one with ‘awareness from the inside’ of what one is doing. Indeed, this is exactly what seems to be *missing* in cases of AHS. Such individuals are often unaware of what their anarchic hand is doing *unless they observe it*. This is in stark contrast with the way that we are usually aware of our actions, and may contribute to the sense of alienation that those with AHS have in response to their bodily movements.

In support of this suggestion, consider the following account of the experience of an AHS patient known as ‘JC’:

For example, in one of the testing sessions, [JC] was asked to turn pages of a magazine with his left hand. As he did this (without any difficulty), the examiner lightly touched his right fingers with a pen. The right [afflicted] hand reached towards and persisted in following the pen continuously as it was slowly moved away from the hand [...]. This reaching continued until the limb was a foot above the table. JC was

unaware of his arm moving on that side (Biran, Giovannetti, Buxbaum, & Chatterje 2006, 567).

In light of this and similar cases, it seems likely that AHS patients lack awareness ‘from the inside’ of their bodily behavior. If the present view is correct, then this is precisely the type of awareness that agentive thoughts provide, and that regular beliefs formed on the basis of conscious observation and inference do not.

### **5.3 Agentive Awareness Is Prior to Agentive Thought**

Another worry for this account is that thoughts that one is acting seem to be *based on* agentive awareness rather than constituting it. Here is Bermúdez (2010) making this point (see also Bayne and Pacherie, 2007, pp. 476):

... beliefs about our own agency need to be anchored in something. They are not free-standing and they do not come out of nowhere. Beliefs are formed for reasons and the most plausible candidates for those reasons relate to our experience as agents (pp. 592).

It is worth emphasizing here that on the present view, agentive thoughts are not ‘free-standing’, nor do they ‘arrive out of nowhere’. But they are not based on experiences of agency either. Rather, they are based on the formation of an executive intention, which may involve an experience when it is conscious, but not one that would constitute agentive awareness. Notice also that the present account of agentive awareness does actually not rule out that agentive awareness is based on antecedent agentive experiences. It might simply be that we sometimes form further beliefs about our own agency on the basis of agentive

experiences constituted by agentic thoughts, much in the same way that we sometimes form perceptual beliefs on the basis of our perceptual experiences.

#### **5.4 Agentic Thoughts Are Too Intellectually Demanding**

Another complaint one might have about the account I have presented is that it is guilty of ‘over-intellectualizing’ agentic awareness. The type of first-person thought that I am positing to explain agentic awareness requires the possession and deployment of sophisticated concepts, for example, the essential indexical and a battery of action concepts, such that creatures with little to no linguistic competence would arguably not be able to form them. But if so, then this may rule out infants and nonhuman animals from being capable of having agentic awareness. When they engage in action, however, they do not behave in ways that indicate alienation from what they are doing. So it would seem that the present account delivers the wrong verdict here.

While perhaps intuitively compelling, this objection relies on a pervasive yet faulty inference. The inference goes from the claim that there is an *absence* of alienation from a certain behavior to the claim that there is a *presence* of agentic awareness for that behavior. But this is not a valid step. The fact that infants and nonhuman animals typically do not react in a surprised or distressed manner when they engage in voluntary behavior, as people with AHS sometimes react to the movements of their affected limb, does not give us reason to conclude that they are agentially aware in those cases—they may simply be engaged in behavior without any distinctive awareness of what they are doing. Without this erroneous inference, we do not have any independent reason to favour the view that nonhuman animals or infants are capable of the type of self-awareness that has been our target phenomenon here, thus leaving the cognitive account untouched on this front.



## **6 Conclusion**

Agentive awareness is commonly thought to be the product of the sub-personal states and mechanisms of the CM. In this paper, I have presented and defended an alternative account, on which it is instead constituted by thoughts that are based on executive intentions. My account thus places agentive awareness squarely in the domain of nonsensory, personal-level states and processes, and gives shape to a relatively overlooked option for how this type of awareness can be best understood, and one that, in my view, deserves further exploration.

*Department of Philosophy and Institute of Cognitive Science*

*Carleton University*

## **Acknowledgements**

For helpful comments and discussion, I am grateful to the anonymous referees for this journal, Jacob Berger, Andrew Brook, Uriah Kriegel, Elisabeth Pacherie, Jake Quilty-Dunn, David Rosenthal, Joshua Shepherd, members of the Oxford Mind Workshop, and audiences at the Institut Jean Nicod PACS group, the Graduate Center, CUNY philosophy colloquium series, and the Carleton philosophy colloquium series.

## References

- Anscombe, G. E. M. 1966. *Intention* (2nd ed.). Cambridge, MA: Harvard University Press.
- Archer, A. 2015. Reconceiving direction of fit. *Thought: A Journal of Philosophy*, 4(3), 171-180.
- Assal, F., Schwartz, S., and Vuilleumier, P. 2007. Moving with or without will: functional neural correlates of alien hand syndrome. *Annals of Neurology*, 62(3), 301-306.
- Bayne, T. 2008. The phenomenology of agency. *Philosophy Compass*, 3(1), 182-202.
- Bayne, T. 2011a. Agentic experiences as pushmi-pullyu representations. In A. A. B. J. H. Aguilar, K. Frankish (eds.), *New Waves in Philosophy of Action* (pp. 219 - 236). New York, NY: Palgrave Macmillan.
- Bayne, T. 2011b. The sense of agency. In F. Macpherson (ed.), *The Senses* (pp. 490 - 524). Oxford: Oxford University Press.
- Bayne, T. and Levy, N. 2006. The feeling of doing: Deconstructing the phenomenology of agency. In W. P. N. Sebanz (ed.), *Disorders of Volition* (pp. 49 - 68). Cambridge, MA: MIT Press.
- Bayne, T. and Montague, M. 2012. *Cognitive Phenomenology*. Oxford, UK: Oxford University Press.
- Bayne, T. and Pacherie, E. 2007. Narrators and comparators: the architecture of agentic self-awareness. *Synthese*, 159, 475-491.
- Banks, G., Short, P., Martinez, J., Latchaw, R., Ratcliff, G., and Boller, F. 1989. The alien hand syndrome: Clinical and postmortem findings. *Archives of Neurology*, 46: 456-459.
- Berti, A. and Pia, L. 2006. Understanding motor awareness through normal and pathological behavior. *Current Directions in Psychological Science*, 15(5), 245-250.
- Biran, I., Giovannetti, T., Buxbaum, L., and Chatterjee, A. 2006. The alien hand syndrome: What makes the alien hand alien? *Cognitive Neuropsychology*, 23(4), 563-582.

- Blakemore, S-J., Wolpert, D., and Frith, C. 2002. Abnormalities in the awareness of action. *Trends in Cognitive Sciences*, 6(6), 237-242.
- Brand, M. 1984. *Intending and Acting*. Cambridge, MA: MIT Press.
- Bratman, M. 1987. *Intention, Plans, and Practical Reason*. Cambridge, MA: Cambridge University Press.
- Burge, T. 2003. Perceptual entitlement. *Philosophy and Phenomenological Research*, 67(3), 503-548.
- Carruthers, G. The case for the comparator model as an explanation for the sense of agency and its breakdowns. *Consciousness and Cognition*, 21, 3-45.
- Chassagnon, S., Minotti, L., Kremer, S., Hoffmann, D., and Kahane, P. 2008. Somatosensory, motor, and reaching/grasping responses to direct electrical stimulation of the human cingulate motor areas. *Journal of Neurosurgery*, 109(4), 593-604.
- Cole, J. 1995. *Pride and a Daily Marathon*. Cambridge, MA: MIT Press.
- Cole, J., Gallagher, S., & McNeill, D. 2002. Gesture following deafferentation: A phenomenologically informed experimental study. *Phenomenology and the Cognitive Sciences*, 1, 49 - 67.
- Davidson, D. 2001. Actions, Reasons, and Causes. In *Essays on Actions and Events* (pp. 43 - 61). Oxford, UK: Oxford University Press. (Original published in 1963.)
- Della Sala, S. 2005. The anarchic hand. *The Psychologist*, 18(10), 606 - 609.
- Desantis, A., Weiss, C., Schutz-Bosbach, S., and Waszak, F. 2012. Believing and perceiving: authorship belief modulates sensory attenuation. *PLoS One*, 7(5), e37959.
- Desmurget, M., Reilly, K.T., Richard, N., Szathmari, A., Mottolese, C., and Sirigu, A. 2009. Movement intention after parietal cortex stimulation in humans. *Science*, 324, 811-813.

- Falvey, K. 2000. Knowledge in intention. *Philosophical Studies*, 99, 21-44.
- Fotopoulou, A., Tsakiris, M., Haggard, P., Vagopoulou, A., Rudd, A., and Kopelman, M. 2008. The role of motor intention in motor awareness: an experimental study on anosognosia for hemiplegia. *Brain*, 131(Pt 12), 3432-3442.
- Fourneret, P., Paillard, J., Lamarre, Y., Cole, J., and Jeannerod, M. 2002. Lack of conscious recognition of one's own actions in a haptically deafferented patient. *Neuroreport*, 13(4), 541-547.
- Fried, I., Katz, A., McCarthy, G., Sass, K. J., Williamson, P., Spencer, S. S., and Spencer, D. D. 1991. Functional organization of human supplementary motor cortex studied by electrical stimulation. *Journal of Neuroscience*, 11(11), 3656-3666.
- Frith, C. 2005. The self in action: lessons from delusions of control. *Consciousness and Cognition*, 14(4), 752-770.
- Frith, C. 2007. *Making Up the Mind: How the Brain Creates Our Mental World*. Oxford, UK: Blackwell Publishing.
- Frith, C. 2012. Explaining delusions of control: The comparator model 20 years on. *Consciousness and Cognition*, 21(1), 52 - 54.
- Frith, C. D., Blakemore, S., and Wolpert, D. M. 2000. Explaining the symptoms of schizophrenia: abnormalities in the awareness of action. *Brain Research: Brain Research Review*, 31(2-3), 357-363.
- Gallagher, S. 2000. Philosophical conceptions of the self: implications for cognitive science. *Trends in Cognitive Sciences*, 4(1), 14-21.
- Gallagher, S. 2007. The natural philosophy of agency. *Philosophy Compass*, 2(2), 347-357.
- Garbarini, F., Rabuffetti, M., Piedimonte, A., Pia, L., Ferrarin, M., Frassinetti, F., Gindri, P., Cantagallo, A., Driver, J., and Berti, A. 2012. 'Moving' a paralysed hand:

- Bimanual coupling effect in patients with anosognosia for hemiplegia, *Brain*, 135, 1486-1497.
- Gennaro, R. 1996. *Consciousness and Self-Consciousness*. Amsterdam: John Benjamins.
- Gennaro, R. 2012. *The Consciousness Paradox: Consciousness, Concepts, and Higher-Order Thoughts*. Cambridge, MA: MIT Press.
- Ginet, C. 1990. *On Action*. Cambridge: Cambridge University Press.
- Giovannetti, T., Buxbaum, L. J., Biran, I., and Chatterjee, A. 2005. Reduced endogenous control in alien hand syndrome: evidence from naturalistic action. *Neuropsychologia*, 43(1): 75-88.
- Guillaud, E., Simoneau, M., and Blouin, J. 2011. Prediction of the body rotation-induced torques on the arm during reaching movements: evidence from a proprioceptively deafferented subject. *Neuropsychologia*, 49(7), 2055-2059.
- Haggard, P., Clark, S., and Kalogeras, J. 2002. Voluntary action and conscious awareness. *Nature Neuroscience*, 5(4), 382-385.
- Haggard, P. and Cole, J. 2007. Intention, attention and the temporal experience of action. *Consciousness and Cognition*, 16(2), 211-220.
- Haggard, P. and Eimer, M. 1999. On the relation between brain potentials and the awareness of voluntary movements. *Experimental Brain Research*, 126(1), 128-133.
- Haggard, P. and Johnson, H. 2003. Experiences of voluntary action. *Journal of Consciousness Studies*, 10(9 - 10), 72 - 84.
- Haggard, P., Newman, C., and Magno, E. 1999. On the perceived time of voluntary actions. *British Journal of Psychology*, 90 (Pt 2), 291-303.

- Horgan, T., Tienson, J., and George, G. 2003. The phenomenology of first-person agency. In S. Walter & H. Heinz-Dieter (eds.), *Physicalism and Mental Causation* (pp. 323 - 340). Imprint Academic.
- Horgan, T. 2012. From agential phenomenology to cognitive phenomenology: A guide for the perplexed. In T. Bayne & M. Montague (eds.), *Cognitive Phenomenology* (pp. 57-78). New York: Oxford University Press.
- Humberstone, I. L. 1992. Direction of fit. *Mind*, 101(401). 59 - 83.
- Jeannerod, M. 2009. The sense of agency and its disturbances in schizophrenia: a reappraisal. *Experimental Brain Research*, 192, 527-532.
- Jenkinson, P. M. and Fotopoulou, A. 2010. Motor awareness in anosognosia for hemiplegia: experiments at last! *Experimental Brain Research*, 204(3), 295-304.
- Koriat, A. The feeling of knowing: some metatheoretical implications for consciousness and control. *Consciousness and Cognition*, 9(2), 149-171.
- Kriegel, U. 2015. *The Varieties of Consciousness*. New York: Oxford University Press.
- Lafargue, G., and Franck, N. 2009. Effort awareness and sense of volition in schizophrenia. *Consciousness and Cognition*, 18(1), 277-289.
- Lau, H. and Brown, R. forthcoming. The emperor's new phenomenology? The empirical case for conscious experience without first-order representations. In A. Pautz & D. Stoljar (eds.), *Themes from Block*. Cambridge, MA: MIT Press.
- Lau, H. C., Rogers, R. D., Haggard, P., and Passingham, R. E. 2004. Attention to intention. *Science*, 303, 1208-1210.
- Lau, H. C., Rogers, R. D., and Passingham, R. E. 2006. On measuring the perceived onsets of spontaneous actions. *Journal of Neuroscience*, 26(27), 7265-7271.

- Lau, H. C., Rogers, R. D., and Passingham, R. E. 2007. Manipulating the experienced onset of intention after action execution. *Journal of Cognitive Neuroscience*, 19(1), 81-90.
- Lau, H. and Rosenthal, D. 2011. Empirical support for higher-order theories of conscious awareness. *Trends in Cognitive Sciences*, 15(8), 365-373.
- Libet, B., Gleason, C. A., Wright, E. W., and Pearl, D. K. 1983. Time of conscious intention to act in relation to onset of cerebral activity (readiness-potential). The unconscious initiation of a freely voluntary act. *Brain*, 106 (Pt 3), 623-642.
- Mandelbaum, E. forthcoming. Attitude, inference, association: On the propositional structure of implicit bias, *Noûs*, 49(1).
- Mandik, P. 2010. Control consciousness. *Topics in Cognitive Science*, 2, 643-657.
- Marcel, A. 2003. The sense of agency: Awareness and ownership of action. In J. Roessler & N. Eilan (Eds.), *Agency and Self-Awareness: Issues in Philosophy and Psychology* (pp. 48-93). Oxford: Oxford University Press.
- Marchetti, C., and Della Sala, S. 1998. Disentangling the alien and anarchic hand. *Cognitive Neuropsychiatry*, 3(3), 191-207.
- McDowell, J. 2010. What is the content of an intention in action? *Ratio*, 23, 415 - 432.
- Mele, A. R. 1992. *Springs of Action*. New York: Oxford University Press.
- Mele, A. R. 2009. *Effective Intentions: The Power of Conscious Will*. Oxford, UK: Oxford University Press.
- Mellors, C. S. 1970. First-rank symptoms of schizophrenia. *British Journal of Psychiatry*, 117, 15-23.
- Metcalf, J., Van Snellenberg, J. X., DeRosse, P., Balsam, P., and Malhotra, A. K. 2012. Judgments of agency in schizophrenia: an impairment in auto-noetic metacognition.

- Philosophical Transactions of the Royal Society of London B Biological Sciences*, 367(1594), 1391-1400.
- Metzinger, T. 2006. Conscious volition and mental Representation: Toward a more fine-grained analysis. In N. Sebanz & W. Prinz (eds.), *Disorders of Volition* (pp. 19-49). Cambridge, MA: MIT Press.
- Millikan, R.G. 1996. Pushmi-pullyu representations. In J. Tomberlin (ed.) *Philosophical Perspectives IX*, 185-200. Reprinted in *Mind and Morals*, ed. L. May & M. Friedman (Eds). Cambridge, MA: MIT Press, 145-61.
- Mylopoulos, M. 2012. Evaluating the case for the low-level approach to agentic awareness. *Philosophical Topics*, 40(2), 103–127.
- Mylopoulos, M. I. 2015. Agentic awareness is not sensory awareness. *Philosophical Studies*, 172(3), 761-780.
- Mylopoulos, M. 2015. Action consciousness and pathologies of agency. In R. J. Gennaro (ed.), *Disturbed Consciousness: New Essays on Psychopathology and Theories of Consciousness*. Cambridge, MA: MIT Press.
- Pacherie, E. 2015. Time to act: The dynamics of agentic experiences. In B. Eitam & P. Haggard (eds.), *The Sense of Agency*. Oxford, UK: Oxford University Press.
- Pacherie, E. 2008. The phenomenology of action: A conceptual framework. *Cognition*, 107, 179-217.
- Pacherie, E., and Haggard, P. 2010. What are intentions? In L. Nadel & W. Sinnott-Armstrong (eds.), *Benjamin Libet and Agency* (pp. 70-84). Oxford, UK: Oxford University Press.
- Peacocke, C. 2007. Mental action and self-awareness (I). In B.P. McLaughlin and J.



- Cohen (eds.), *Contemporary Debates in Philosophy of Mind*, Oxford: Blackwell: 358-376.
- Paul, S. 2009. How we know what we're doing. *Philosophers' Imprint*, 9(11), 1-24.
- Pitt, D. 2004. The phenomenology of cognition or what is it like to think that p? *Philosophy and Phenomenological Research*, 69(1), 1-36.
- Prinz, J. J. 2007. All consciousness is perceptual. In B. P. McLaughlin & J. Cohen (eds.), *Contemporary Debates in Philosophy of Mind*. Malden, MA: Blackwell Publishing Ltd.
- Prinz, J. J. 2012. *The conscious brain: How attention engenders experience*. New York, NY: Oxford University Press.
- Proust, J. 2013. *The Philosophy of Metacognition: Mental Agency and Self-Awareness*. Oxford, UK: Oxford University Press.
- Ramachandran, V.S. and Hirstein, W. 1998. The perception of phantom limbs. The D. O.Hebb lecture. *Brain*, 121 (Pt 9), 1603–1630.
- Rosenthal, D. M. 1986. Two concepts of consciousness. *Philosophical Studies*, 49: 329-359.
- Rosenthal, D. M. 2005. *Consciousness and Mind*. New York: Oxford University Press.
- Rosenthal, D. 2011. Exaggerated reports: reply to Block. *Analysis*, 71(3), 431-437.
- Schneider, K. 1957. Primary & secondary symptoms in schizophrenia. *Fortschritte der Neurologie – Psychiatrie*, 25(9), 487–490.
- Sarlegna, F. R., and Bernier, P. M. 2010. On the link between sensorimotor adaptation and sensory recalibration. *Journal of Neuroscience*, 30(35), 11555-11557.
- Sarlegna, F. R., Gauthier, G. M., Bourdin, C., Vercher, J. L., and Blouin, J. 2006. Internally driven control of reaching movements: a study on a proprioceptively deafferented subject. *Brain Research Bulletin*, 69(4), 404-415.

- Schurger, A., Mylopoulos, M., Rosenthal, D. 2015. Neural antecedents of spontaneous voluntary movement: A new perspective. *Trends in Cognitive Sciences*, 20(2), 77-79.
- Schurger, A., Sitt, J. D., and Dehaene, S. 2012. An accumulator model for spontaneous neural activity prior to self-initiated movement. *Proceedings of the National Academy of Science U S A*, 109(42), E2904-2913.
- Schwenkler, J. 2012. Non-observational knowledge of action. *Philosophy Compass*, 7(10), 731-740.
- Searle, J. R. 1979. The intentionality of intention and action. *Inquiry*, 22(1-4), 253-280.
- Searle, J. R. 1983. *Intentionality: An Essay in the Philosophy of Mind*. Cambridge, MA: Cambridge University Press.
- Sellars, W. 1976. Volitions re-affirmed. In D. W. M. Brand (ed.), *Action Theory* (pp. 47 - 66). Hingham, MA: D. Reidel Publishing Company, Inc.
- Setiya, K. 2007. *Reasons Without Rationalism*. Princeton: Princeton University Press.
- Shepherd, J. 2016. Conscious action/zombie action. *Noûs*, 50(2), 419-444.
- Siegel, S. 2010. *The Contents of Visual Experience*. Oxford, UK: Oxford University Press.
- Smith, M. 1987. The Humean theory of motivation. *Mind*, 96(381), 36-61.
- Spehn, M. K. and Reder, L. M. (2000). The unconscious feeling of knowing: a commentary on Koriat's
- Spence, S. A., Brooks, D. J., Hirsch, S. R., Liddle, P. F., Meehan, J., & Grasby, P. M. 1997. A PET study of voluntary movement in schizophrenic patients experiencing passivity phenomena (delusions of alien control). *Brain*, 120 (Pt 11), 1997-2011.
- Synofzik, M., Vosgerau, G., and Newen, A. 2008. Beyond the comparator model: A multifactorial two-step account of agency. *Consciousness and Cognition*, 17(1), 219 - 239.
- Tsakiris, M., Shütz-Bosbach, S., Gallagher, S. 2007. On agency and body-ownership:

- Phenomenological and neurocognitive reflections. *Consciousness and Cognition*, 16, 645-660.
- Velleman, J. D. 1989. *Practical Reflection*. Princeton, NJ: Princeton University Press.
- de Vignemont, F. and Fournieret, P. 2004. The sense of agency: A philosophical and empirical review of the 'Who' system. *Consciousness and Cognition*, 13, 1 - 19.
- Vocat, R., Saj, A., and Vuilleumier, P. 2013. The riddle of anosognosia: Does unawareness of hemiplegia involve a failure to update beliefs? *Cortex*, 49(7), 1771-1781.
- Vuilleumier, P. 2004. Anosognosia: the neurology of beliefs and uncertainties. *Cortex*, 40(1), 9-17.
- Wegner, D. 2002. *The Illusion of Conscious Will*. Cambridge, MA: Bradford Books.
- Wolpert, D. M. 1997. Computational approaches to motor control. *Trends in Cognitive Science*, 1(6), 209-216.
- Wolpert, D. M., Ghahramani, Z., & Jordan, M. I. 1995. An internal model for sensorimotor integration. *Science*, 269(5232), 1880-1882.