***Penultimate Draft: Published in J. Weisberg & D. Rosenthal (eds.), *Consciousness* (pp. 280-293). Hoboken, NJ: Wiley-Blackwell, 2022. Please cite published version only.**

The Proprietary Nature of Agentive Experience

Myrto Mylopoulos

(Carleton University)

1. Introduction

The main contention of this paper is that, just as there is something it is like to smell a rose, taste chocolate, and hear a siren, there is something it is like to perform an action. In other words, I will argue that we ought to recognize, alongside these other familiar forms of phenomenology, a distinctive *phenomenology of agency*.[1] My claim is not simply that there is some subjective experience that attaches to the performance of actions. No one disputes that action is typically accompanied by a range of kinesthetic and visual experiences, and often preceded by conscious thoughts about what to do. The position I wish to defend, rather, is that there is a *proprietary* phenomenology or subjective character associated with action that "goes beyond" (Horgan et al. 2003, p. 323) these familiar types of sensory and cognitive experience (for other treatments of this topic, see Horgan 2007; Bayne and Levy 2006; Bayne 2008; Pacherie 2008; Kriegel 2015; Shepherd 2017).

What is it for some type of experience to be proprietary in the sense that I have in mind? At the very least, it must not be such that it can simply be subsumed under, or viewed as a determinate of (see Kriegel 2015 for this framing), another experience type with which we are familiar. The experience of seeing color, for instance, does not meet this condition, since it can be subsumed under, or viewed as a determinate of, visual experience. Visual experience, by contrast, cannot be so treated. So it meets the condition in question. It is in this sense that I take agentive experience to be proprietary: it cannot be subsumed under, or viewed as a determinate of, familiar types of cognitive or sensory experience.

What is it for some type of proprietary experience to be *agentive*? The short answer is that it must be distinctive of action. Clearly it will not do for the experience in question to

---

[1] This has also been referred to as the "sense of agency" (e.g., Marcel 2003) and "agentive awareness" (e.g., Mylopoulos 2015, 2017).

accompany actions only rarely, or to frequently accompany cases of non-action, such as passive movement or involuntary reflex. We are looking for something like a subjective signature of action performance. Still, we should not demand that agentive experience accompanies all and only actions. For instance, one might dream that one is flying, and have an accompanying agentive experience, though of course one is not performing any action in this case. Or one might sleepwalk without any agentive experience, though one is arguably engaged in an action here. To leave room for such cases in which the mechanisms by way of which agentive experience is generated are sometimes faulty, we should say that proprietary agentive experience is what is *typical* of occurrences of action. What we want, then, is an experience type that is well-correlated with action: it is typically present when action is present, and typically absent when action is absent. Taking all this on board, we have the following characterization: An experience is proprietarily agentive just in case it cannot be subsumed under, or treated as a determinate of, familiar types of cognitive or sensory experience, and it is well-correlated with action performance.

Now it might be thought that the present undertaking is unnecessary, at best, and misguided, at worst. After all, why bother to *argue* for the existence of proprietary agentive experience if we can just use introspection to establish it? But there are familiar problems with appealing solely to introspection in order to determine the character of one's conscious experiences. If we look to introspection alone in order to ascertain whether or not there are proprietary agentive experiences, then there is no progress to be made in the face of disagreement.

Suppose that someone sincerely maintains that when they "look inside" as they are acting, they find no agentive experiences of which to speak. As their body moves, it seems to them that they have all sorts of conscious proprioceptive, visual, and kinesthetic experiences, and even some conscious thoughts about what they are doing, but no proprietary *agentive* experiences that they can discern. Suppose now, that a second someone, feels quite differently. When *they* introspect as they act, they feel that they have a clear conscious experience of being the agent of their action that goes over and above any sensory or cognitive phenomenology. What can be said here to settle this disagreement? Since each side has only their introspective data to go on, there is no independent tribunal to which to appeal in order to resolve the conflict.

This means that we need to construct a theoretical case in favor of positing proprietary agentive experiences. This is what I attempt to do in the remainder of this paper.

2. Proprietary Agentive Experience and Two Types of Skepticism

Consider the following vignette:

This morning I was awakened by the sound of someone practising the violin. I dozed a bit, then got up, washed, shaved, dressed, and went downstairs, turning off the light in the hall as I passed. I poured myself some coffee, stumbling on the edge of the dining room rug, and spilled my coffee fumbling for the *New York Times*. (Davidson 2001, p. 43)

<P>Within this familiar scene, there are things that happen to the agent—being awakened, dozing, stumbling, spilling, and fumbling—and there are things that the agent does, that is, actions the agent performs—getting up, washing, shaving, dressing, going downstairs, turning off the light, and pouring coffee. Our starting point in this discussion is to say that there is *some* phenomenal or subjectively experienced contrast to be drawn between all the episodes in the former category and those in the latter category. The subjective experiences accompanying the events in the former category are different than those accompanying the events in the latter category. What's more, the contrast in question cannot be entirely explained by the varying nature of the activities and happenings in each category. It has also to do with the fact that the latter are actions and the former are not. For we could make the same point using one and the same bodily movement, in one case actively generated and in another passively generated. There is a phenomenal difference between, for instance, raising one's arm of one's own volition and having someone else raise one's arm passively, but producing just the same motion.

I shall defend the view that the phenomenal contrast between the experiences involved in action versus passive movement is best explained by appeal to proprietary agentive experience that accompanies the former, but not the latter. In order to do so,

however, we must address two kinds of skepticism about the source of the agreed upon phenomenal contrast. These I refer to as *strong* and *moderate* skepticism.[2]

The strong skeptic opts to explain the phenomenal contrast just described by maintaining that action performance *lacks* a subjective quality that passive movement *has*, and that there is no positive subjective experience to assign to action itself. On this view, what is doing all the work in driving the contrast is, rather, the special experiential quality that accompanies our passive or accidental movements. As Prinz (2012) puts the view, "[t]here is no phenomenology of being a controller, but there is a phenomenology of being controlled" (p. 239).

The moderate skeptic, by contrast, adopts a slightly different tactic. Unlike the strong skeptic, they allow that there is agentive experience that accounts for the phenomenal contrast, but they deny that this type of experience is proprietary. Rather, they hold that the agentive experience can be reduced to "intentions, perceptual expectations, and perceptual experience" (Grünbaum 2015, p. 3318), that is, familiar types of cognitive and sensory phenomenology associated with action production, and thus fails to be proprietary in the sense that we have laid out.

Having identified the two sources of skepticism that must be addressed in order to establish the existence of proprietary agentive experiences, I turn now to one existing strategy for doing so. While ultimately I do not think the strategy works, I will use it as a springboard for the strategy that I pursue in the rest of the paper.

3. The Robust Conceivability Strategy

Horgan (2012) attempts to establish the existence of agentive phenomenology in the face of skepticism by appealing to the robust conceivability of certain kinds of zombie scenarios. Robustness here involves "inter alia, the persistence of conceivability under inclusion of

---

[2] I set aside what might be called the *extreme* skeptical view that denies that there is any phenomenal contrast in these cases. I also set aside what might be labeled the *weak* skeptical view that, while readily acknowledging that there are phenomenal differences between the conscious experience of acting and the conscious experience of being moved passively, maintains that, insofar as agentive phenomenology is supposed to involve a *type* of subjective experience that typically accompanies actions and fails to accompany passive movements, it does not exist. Rather, each *token* action has some agentive phenomenal quality that distinguishes it from passive movement, but this same quality does not attach to *all* action (see Bermúdez 2010).

arbitrarily greater detail and specificity" (p. 60). In order to get clear on Horgan's strategy, it will be useful to briefly review the nature of a philosophical zombie. A *complete* zombie is a creature that is a physical and functional duplicate of a human being, yet has no conscious experiences whatsoever. A *partial* zombie is a creature that is a physical and functional duplicate of a human, and has *some* conscious experiences, but is lacking certain other types of conscious experience that it is agreed that humans typically have. Horgan's strategy is to present a series of thought experiments involving "robustly conceivable creatures who (i) are complete functional duplicates of ordinary humans, (ii) have the same sensory phenomenology as ordinary humans, … but (iii) are partial zombies nonetheless" (p. 61). On Horgan's view, what these creatures are missing are certain kinds of *non-sensory* experience "that are present in the mental lives of ordinary human beings" (2012, p. 61). His ultimate aim is to show that this verdict applies to cognitive phenomenology of the sort that accompanies conscious thought, but as a stepping stone to this conclusion, he starts by attempting to show that it applies to agentive experiences.

Let us start with a regular human being who we will call Simone1. First, Horgan asks us to conceive of a creature that is a duplicate of Simone1 from the third-person perspective. We can call her Simone2. Simone2 is to be conceived of as having the exact same functional organization and physical implementation of her functional states as Simone1. Next, Horgan asks us to conceive of Simone2 largely from the first-person perspective, with all sensory experiences intact, but "no experiences as-of certain bodily movements emanating from h[er]self as their source" (2012, p. 67). Finally, we are asked to conceive of Simone2 as having all the Stage 1 and Stage 2 features that we attributed to her, combining them to get her full phenomenological profile.

Horgan maintains that a creature just like Simone2 is easily and robustly conceivable, and that to deny this would be "a very large bullet to bite" (2012, p. 69). He goes on to claim that, since Simone1 and Simone2 are physico-functional duplicates, the only differences between the conceived of creatures must be differences in their conscious experiences—in particular, Simone1 has and Simone2 lacks agentive experiences of her movements. Thus, according to Horgan, "the robust conceivability of certain kinds of zombie scenarios can serve as a criterion for the existence of certain kinds of phenomenal character" (p. 62).

Unfortunately, Horgan's strategy fails to get us anywhere in dealing with either of the skeptics we have just introduced. There are a number of serious difficulties it faces. For one,

the way the scenario in the second step of the strategy is described is problematic. The key instruction is to conceive of Simone2 as lacking "experiences as-of certain bodily movements emanating from h[er]self as their source" (p. 67). But what exactly is one being instructed to do here? Horgan offers some elaboration, however, in doing so he ends up corrupting the results of the conceivability test, for whatever they might have been worth. For, in expanding on what it means to conceive of Simone2 as lacking such experiences, he urges that we conceive of her as being such that she "always experiences [her] own bodily movements as just happening, in much the way one experiences one's lower leg extending itself when a doctor taps one's knee with a reflex-testing mallet" (Horgan 2012, p. 67). But this is compatible with the moderate skeptic's view that there are no agentive experiences, just the presence (or absence) of passive experiences, and this is what allows for the relevant conceivability. In order to comply with the instructions, one need not conceive of the absence of any proprietary type of agentive experience, just the pervasive presence of what it feels like to be passively moved. So this will not help to dispel the worries of our moderate skeptic.

In a similar vein, note that one would be able to comply with the instructions even if one did not take agentive experiences to exist. It would be as though one were asked to conceive of a partial-zombie creature that is a functional duplicate of another creature, but without the conscious experience of their hair growing. If one does not think such an experience type exists, one will just conceive of this creature as being the same as their duplicate, but not by "subtracting" anything from their experience, since there is, one believes, nothing to subtract. The same might be true of the conceivability of Horgan's partial zombie.

Finally, there is a more theoretical worry lurking here: it may be that one cannot conceive of the relevant partial zombie scenario because one cannot conceive, more broadly, of the schism between physical-functional properties and phenomenal properties that this requires. Remember that Simone2 is supposed to be *exactly* alike in her physical and functional organization as Simone1. This means that she carries herself around the world, engaging in well-controlled actions, and even talking about things she has done, without reporting anything amiss. One might not be able to conceive of someone who, from the outside, appears to be an intact, successful agent in terms of their functioning and yet experiences all of their movements as reflex-like. This has no bearing, however, on the

existence of agentive experience, but rather the relationship between phenomenal properties and physical-functional properties.

Ultimately, I think a less roundabout approach is needed than that offered by the robust conceivability strategy. I do think Horgan is on the right track, however, in attempting to narrow in on a clearer description of what exactly agentive experience amounts to—and in particular his description of it as involving a sense of oneself as the source of an action. There will be more on this below.

In the meantime, I propose the following approach: the two forms of skepticism we have laid out (strong and moderate) earn their keep insofar as they offer competing explanations of the accepted phenomenal contrast between performing an action and passive movement—explanations that do not appeal to proprietary agentive experiences. Not only this, but I take them to offer the *only* competing explanations. Granting that there is a subjective difference between action and passive movement, there must be some explanation of it, and I do not see credible alternatives other than allowing that the contrast is proprietary, but denying that it is agentive, as the strong skeptic maintains, or allowing that it is agentive, but denying that it is proprietary, as the moderate skeptic does.

In order to defeat this skepticism, then we need to show that these explanations cannot actually account for the phenomenal contrast between acting and passively moving. And part of this will involve describing in fuller detail what agentive experiences amount to. If we succeed, then we will be left with the one remaining explanation: there is a proprietary phenomenology of agency. Let's take on the strong skeptic first.

4. Confronting the Strong Skeptic

The strong skeptic maintains that there is a difference between the experience of acting and that of passively moving only insofar as the experience of passively moving has its own proprietary phenomenology, and not insofar as the experience of acting does.

In responding to this type of skepticism, it will be useful to have a clear example before us of a case for which one might be tempted to adopt it. Consider the neuropsychological condition known as 'anarchic hand syndrome' (AHS), which results from damage to the supplementary motor area (SMA) or anterior corpus callosum (Della Sala 2005). In this condition, individuals execute complex, goal-oriented movements with an upper limb, usually that which is contralesional, all the while denying authorship of those

movements (Marchetti and Della Sala 1998). Individuals afflicted with AHS are unable to inhibit the movements of their anarchic limb, except by indirect methods, such as using their other hand to stop it. They experience significant distress and frustration at their condition (see Biran et al. 2006).

The neurologist Sergio Della Sala (2005) describes a particularly striking episode involving one of his anarchic hand patients:

One evening we took our patient, Mrs. GP, to dinner with her family. We were discussing the implication of her medical condition for her and her relatives, when, out of the blue and much to her dismay, her left hand took some leftover fish-bones and put them into her mouth. A little later, while she was begging it not to embarrass her any more, her mischievous hand grabbed the ice-cream that her brother was licking. Her right hand immediately intervened to put things in place and as a result of the fighting the dessert dropped on the floor. She apologised profusely for this behaviour that she attributed to her hand's disobedience. Indeed she claimed that her hand had a mind of its own and often did whatever "pleased it." (p. 606)

Mrs. GP does not identify as the agent of the movements of her anarchic limb, even attributing to it a mind of its own.[3]

In response to such a case, Bayne (2011) has argued that the denials of Mrs. GP can be explained by appeal, at least in part, to the fact that the agentive experiences that normally accompany action have been replaced by passive experiences of things merely happening to her. He writes,

[3] We could appeal to other cases from abnormal psychology as well: Schizophrenic individuals undergoing delusions of control attribute their actions to external agents (Frith et al 2000). And in what is known as functional movement disorder (FMD), individuals present "unwanted muscle movement, such as tremor or dystonic posturing, which cannot be voluntarily controlled" (Fried et al. 2017, p. 10844) and likewise deny being the authors of these behaviors. Note that these denials do not amount to denials that the movements emanate from one's body (De Vignemont 2007). Rather, they are denials that one is, oneself, their causal source or author.

It seems plausible to appeal to agentive experience—or the lack thereof—in order to *explain* why [anarchic hand patients] deny having performed the anarchic actions. Surely it is the fact that the normal and expected experience of *doing* has been replaced by an experience of *happening* that leads these patients to judge that the action is not theirs. (p. 498, emphasis in original)

Here is precisely where the strong skeptic may protest. For the strong skeptic will argue that we need not view the experience of happening as *replacing* a feeling of doing. Rather, we can simply suppose that Mrs. GP and other anarchic hand patients experience a new feeling of passivity or alienation in relation to the anarchic limb's movements, where *no* distinctive experience of acting was present before, and this is what drives their reports.

What to say to the strong skeptic? The first thing to note is that the strong skeptical stance has its greatest appeal if we are discussing only simple cases of action. But, though one might not know it from looking at much of the philosophical literature, agents engage in all sorts of complex actions that go beyond raising one's arm, or reaching for one's coffee mug. As a tiny sampling, let's remind ourselves that agents do things like: flip omelets, diffuse bombs, take out the garbage, perform heart surgery, get dressed, hail taxis, thread needles (in the sewing sense), thread the needle (in the dancing sense), and so on. The strong skeptic has to say that in none of these cases can we experience ourselves as agents—all we can experience is losing control of such actions. But given the complex nature of these various action types, this starts to look like the wrong verdict.

Even when it comes to relatively more mundane cases—such as reaching for one's coffee mug—we must not lose sight of the fact that these actions, like all actions, are events that unfold over time and involve coordination and interaction among various psychological capacities (e.g., motor planning, memory, attention, perception) as well as token mental representations of different types (e.g., belief, desire, intention, decision, perception, imagery). The simple action of reaching for one's coffee mug, at minimum, involves a proximal intention to reach for the mug, the perception of the mug, the direction of visual attention toward it, and a cascade of motor commands guiding the action to completion with the help of multimodal sensory feedback. Given the many psychological capacities, representational components and processes involved in even mundane exercises of agency,

we should not only expect a distinct type of agentive experience to distinguish between cases of agency and passive movement, but *multiple* types of agentive experience.

Indeed, Mylopoulos and Shepherd (2020) distinguish among six different types of related and sometimes dissociable experiences that accompany action. These include: the experience of purposiveness, the experience of mineness, the experience of control, the experience of action execution, the experience of action perception, and the experience of libertarian freedom (see, e.g., Strawson 2010 for further discussion of this experience type). I will focus here on the experience of mineness and the experience of control, as I take these to be the most central.

Let's look first at the experience of 'mineness,' 'ownership,' or 'authorship' of an action. This is the experience of an action as one's own, or of oneself as performing an action. The description of this experience may seem puzzling on its face. How could one experience one's actions as anything *but* one's own? One way of thinking about this, which will not be the way we have in mind here, concerns what Velleman (1992) has called the *problem of the disappearing agent*. The idea of the disappearing agent been invoked to describe cases where agents do not "participate" in their actions, and thus fail to be present in their production. The paradigm case is that of an unwilling nicotine addict, finding themselves lighting up another cigarette despite having formed the decision not to. The thought is that the agent is "alienated" from their own action, resulting as it does from desires that they do not endorse.

But this is not what is meant here by an experience of ownership over an action. While it may be true that the unwilling addict does not identify with or endorse the desire that drives their smoking behavior, they still experience themselves as the source of the action. As Schlosser (2011) correctly notes, "it is uncontroversial that the addict's desire [to smoke a cigarette] and the resulting behaviour is his *own* in some basic or minimal sense" (p. 25). The agent is aware of their action as their own in a way that can be reasonably contrasted with how Mrs. GP seems to experience her actions. When anarchic hand patients say things like, "Of course I know it's me, it just doesn't *feel* like me" (as reported in Marcel 2003, p. 79), this is the experience they are alluding to missing.

There is another feature of the anarchic hand case that can be contrasted with the regular experience of authorship. Notice first that there are different ways in which one might come to be aware of oneself as doing something. Suppose, for example, that I notice

an annoyed expression on my neighbor's face, leading me to become suddenly aware that I am tapping my pen on the table. My pen-tapping is an action of mine, but in this case, I come to be aware of it in an unusual way, that is, by consciously observing what I am doing after being tipped off by my neighbor's facial expression. This is *not* the way in which we typically come to be aware of our own actions. Danto and Morgenbesser (1963) nicely depict the relevant contrast:

If one day I should notice that my arm was rising and lowering, and then realize that, if I had not noticed, I would not have known it was doing this, it would be for me a terrifying experience, a sign that I had lost contact with part of myself, that my arm had become an alien entity. (p. 441)

When it comes to awareness of our own actions, unlike the awareness we have of the actions of others, the awareness comes about in a subjectively unmediated way: it does not seem to rely on observation or inference. It appears, however, that anarchic hand patients, even though they are able to attribute the behavior of the anarchic limb to themselves, sometimes only gain awareness of what their anarchic limb is doing "from the outside," through conscious observation of its activity.

Consider the case of a patient known as JC, who acquired AHS at the age of fifty-six after suffering a stroke. JC's anarchic limb executes well-coordinated movements at the level of sensorimotor control, such as reaching and grasping for objects, turning taps on and off, pressing buttons, and so on. But JC is not aware of the movements as his own, reporting that the affected hand "[h]as a mind of its own" and "[w]ants to be the boss" (Biran et al. 2006, p. 567).

There is good reason to think that JC may only be aware of his anarchic movements when he is consciously observing them. This seems to be what is happening in the following incident, for example, in which JC, busy with a different task and therefore not observing his "anarchic" limb, is described as being unaware of the limb's behavior:

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 et al. 2006, p. 567)

If AHS patients are, at least sometimes, not aware of the movements of their anarchic limb unless they consciously observe them, then this is clearly at odds with how we regularly come to be aware of ourselves as performing some action. The experience of authorship, then, may be thought of as a form of subjectively unmediated awareness of oneself as acting.

To get a further sense of what the experience of authorship is thought to consist of, consider the description Horgan (2007) offers of what he calls the *experience of self as source*. This seems to be one important aspect of the experience of authorship. Remarking on the experience of raising one's hand and clenching one's fingers, he writes:

You experience your arm, hand, and fingers as being moved *by you yourself*—rather than experiencing their motion either as fortuitously moving just as you want them to move, or passively experiencing them as being caused by your own mental states. You experience the bodily motion as caused by *yourself*. (Horgan 2007, p. 8)

The quotation is useful as it presents us with a more detailed pair of contrast cases that can be used to hone in on the experience of authorship in ways that our starting contrast between moving one's arm and having it passively moved by an external source cannot. In particular, it helps us identify different ways in which one might experience the connection between oneself and one's actions and contrast these with typical experiences of action. First, in typical experience of action, the link between the motion of one's body and what one wants or intends to do is not experienced as "fortuitous." It is not, for instance, like the experience of watching a penalty shootout in a soccer match, wherein the striker kicks the ball past the goalkeeper and into the net in just the way that one wants them to (top right corner and all). When one acts, one is not, as in this case, aware of one's desire for some outcome to obtain and then aware of it obtaining. Rather, there is a clear subjective sense that one's action is caused by a source internal to oneself. And nor, if Horgan is correct, one merely aware of one's desire or intention as being the cause of one's movement, without being aware of oneself as the cause or source. Rather, awareness of oneself as the cause of

the action seems to be at the core of the experience. After all, many times one finds oneself doing something without any awareness of desiring or intending to do it (e.g., flipping a light switch on the way out of a room). But one may still be aware of oneself as doing it.

The experience that Horgan highlights seems to be absent in cases of delusions of control in schizophrenia. Such delusions are a subtype of what researchers working on this condition call passivity experiences (Frith et al. 2000). In general, passivity experiences involve the attribution of an action, mental event, or sensory experience to an external source. They include a range of phenomena, such as thought insertion ("Thoughts come into my mind from outer space."), somatic passivity experiences ("I have tingling feelings in my legs caused by electric currents from an alternator."), and made emotions ("It puts feelings into me: joy, happiness, embarrassment, depression. It just puts it in and I feel the glow spread over me.") (all quotations from Frith et al. 2000, p. 358).

In delusions of control, schizophrenic individuals experience their actions passively—attributing them to alien sources. Their reports are rather striking. For instance, one person suffering from such delusions reported that "[m]y grandfather hypnotized me and now he moves my foot up and down" (Frith et al. 2000). Another individual insisted that, "the force moved my lips. I began to speak. The words were made for me" (Mellors 1970). Yet another patient, having just made an arm movement, complained about the movement that: "I felt like an automaton, guided by a female spirit who had entered me during it" (Spence et al. 1997). Experiences of authorship seem to be absent in each of these cases. The striking nature of these reports serves to underscore the regularity with which we experience ourselves as the source of our actions.

In addition to the experience of authorship, a second aspect of agentive experience that we can highlight and describe is the experience of controlling one's action or the "sense of control" (e.g., Haggard and Chambon 2012, p. R390). When one engages in action, one often has a sense of controlling it to various degrees, and in addition a sense of being able to modulate this degree of control by way of exercising various capacities. Imagine, for instance, that you are rushing to complete a task in the kitchen, say chopping the vegetables before the water in the pot boils. A familiar experience in this scenario is the feeling that, as you hurriedly chop, you are going too quickly and at risk of cutting yourself. Here your sense of not having enough control over what you are doing might result in your slowing down and attending more carefully to your chopping in order to avoid losing a finger. In other

cases, by contrast, you might feel that you are fully in control of what you're doing, and thereby direct your focus and attention toward other simultaneous tasks. Imagine, for instance, that you are a veteran driver, driving your car along a familiar route under comfortable road conditions. Here you might experience a high degree of control over your actions and so choose to engage your passenger in conversation or flip through stations on the radio.

The experience of control is primarily concerned with the evaluation of an action's success as it unfolds, whereas the experience of authorship is a more general experience of being the causal source of one's action. To appreciate a case in which the two can be clearly distinguished, imagine writing out your name with your non-dominant hand and then doing so with your dominant hand. You will have an experience of authorship for both actions, but an experience of greater control for the action of the dominant hand.

There may also be cases where the experience of authorship is modulated, but the experience of control is not. Consider, as candidates, cases of flow experience that sometimes accompany expert performance (Csikszentmihalyi 2000). Describing such experiences as cases of "fully absorbed coping," Dreyfus (2007) writes,

In fully absorbed coping, there is no immersed ego, not even an implicit one. The coper does not need to be aware of himself even in some minimal way but only needs to be capable of entering a monitoring stance in the brain, which is comparing current performance with how things went in the past, sends an alarm signal that something is going wrong. (pp. 374)

Dreyfus is likely overstating things here in claiming that one is not aware of oneself at all in such cases. But these do seem to be cases in which one's awareness of authorship is at least attenuated, while one's experience of control remains steady. This does not entail, of course, that individuals experience *not* being the authors of their actions in such cases, or that they deny being the authors of their actions. The thought, rather, is that the experience of authoring these actions is dampened or at least modified to some degree relative to cases where one is not immersed in an expert performance (see Dow 2017 for further discussion of the experience of agency during flow).

Taking stock: First, we pointed out the complexity of human agency, both in terms of the range of actions that humans perform and the range of capacities engaged in their performance. Second, we described in some detail two distinct, yet co-occurring aspects of agentive experience: the experience of authorship and the experience of control. How does all this help us to respond to the strong skeptic? It does so by making far less credible the claim that there is *no* positive experiences associated with action that contribute to the phenomenal contrast we began with between raising one's arm and being passively moved. Not only can we clearly describe the character of these experiences, but also cases in which they are modulated or even absent. This goes well beyond what a mere sense of passivity can explain.

It does not yet get us to the point of being able to say that there is a proprietary experience of agency, however. This is because the moderate skeptic can now weigh in: Yes, there are experiences of authorship and control, the moderate skeptic might say, but they are not proprietary. Rather, they may be subsumed under more general types of familiar experience pertaining to action, to wit, cognitive and sensory phenomenology. It is to this form of skepticism that we now turn.

5. Confronting the Moderate Skeptic

The second skeptical stance we are responding to holds that agentive experiences are not proprietary, but rather reducible to co-occurring experiences of more basic, familiar types of action-related experience, such as the cognitive experience of intending or deciding to act, followed by perceptual experiences of bodily movement (perhaps in both visual and kinesthetic modalities).

Shepherd (2017) has labeled this the 'co-consciousness' view of the experience of acting, the thought being that it is characterized by a set of experiences, each of them at the same time conscious, but without any unifying integration among them. If this view is correct, then the experience of agency is no more proprietary than any other experience that involves the co-occurrence of multiple types of experience.

We are familiar with this kind of experience within a single modality. Consider, for example, the experience of watching fireworks, involving as it does an array of distinct colors, patterns of motion, and shape all co-occurring within vision. Or the experience of listening to a symphony orchestra, with the sounds of its string instruments, woodwinds,

brass instruments, and percussion. We are also familiar with this kind of experience across sensory modalities, for example, seeing a flash of lightning followed by the sound of a burst of thunder. The thought is that agentive experiences are akin in that they are co-occurring experiences of intending or deciding to act, followed by sensory experiences from the body and one's environment.

Let's see if we can make good on the moderate skeptic's proposal with respect to the experience of authorship. Can this be subsumed under a familiar experience type? One suggestion that can be found in the literature is that the experience of oneself as performing some action is simply an experience of trying, that is, a case of cognitive phenomenology. Kriegel (2015), takes this (what he calls the experience of "deciding-cum-trying") to be at the core of agentive phenomenology describing it as "a nonsensory analog of innervation (a feeling of a kind of nonsensible current traveling from will to muscle)" (Kriegel 2015, p. 95). Shepherd (2016), too, discusses such experiences, describing them as experiences "as of directing effort towards the satisfaction of an intention" (p. 422) and suggesting further that the neural activity realizing experiences of trying "might simply be the activity that realizes something like a conscious intention at work" (p. 422).

Certainly there is something to the idea that, already within the experience of trying, one has an experience of *oneself* directing effort toward an action outcome. Thus, consider Shepherd's fuller description of this experience:

The directive character of experiences of trying … does not emanate from any bodily location. It is not incorrect to call it an experiential mandate. *But in this case the mandate seems to emanate from the agent. When I have an experience of trying to raise my arm, I have an experience as of mandating that my arm rise.* (Shepherd 2016, p. 421, emphasis added)

If this is correct, then perhaps the experience of authorship, involving as it does the sense of self as source, is simply an experience of trying. And perhaps, going along with Shepherd's (2016) proposal, experiences of trying are simply conscious intentions "at work." This view is tempting, but there is a clear obstacle to this proposal, which is that the experience of authorship involves awareness of oneself as *doing* something, and not merely *trying or intending* to do something.

Kriegel (2015) attempts to offer us a way out of this worry. 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 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. (p. 90)

Does this help? It is true that nothing in the experience of trying indicates the success or failure of the trying. But notice that in the visual case, the hallucinatory awareness and the veridical awareness are still both such that one is aware of oneself as seeing something. So, if the analogy is strictly to hold, both successful tryings and unsuccessful tryings should yield a corresponding awareness of intending to do something, since tryings are just conscious intentions. But this just takes us back to precisely the problem we are attempting to solve. The experience of authorship is awareness of oneself as doing something, not of intending to do something.

In addition, there is ample empirical support for the claim that experiences of trying are subjectively distinguishable from experiences of oneself as acting. Consider evidence from self-paralysis studies, which Shepherd (2016) appeals to in order to support the view that "experiences of trying to move appear to be causally upstream of, and to correlate reliably with, experiences of the body moving" (p. 426). On this view, experiences of trying to move and experiences of moving are clearly distinct. In one such study, participants were administered paralyzing agents that would block feedback from afferent nerves. The authors found that "[a]ttempted voluntary movement of a limb paralysed with intravenous atracurium was accompanied by a marked sense of effort" (Gandevia et al. 1993, p. 85). In particular all three participants in the study "reported strong sensations of effort accompanying attempted movement of the limb, as if trying to move an object of immense weight" (p. 97). These experiences were sometimes, but not always, accompanied by illusory

movements, and in every case clearly distinguishable from the experience of those movements themselves.

For further support that experiences of oneself as acting can be distinguished from the experience of trying, especially if such experiences are construed as being constituted by conscious intentions at work, we can draw on the results of now-classic and well-replicated subjective timing experiments. In such studies, participants are asked to report the time at which they first become aware of a decision to act—the outcome of which is an intention to do so—and in another condition the time at which they first become aware of having acted (see Libet et al. 1983; Haggard and Eimer 1999). Though the temporal window between average reporting times for each event is small, it is not nonexistent. On average, participants report being aware of a decision to act around 200 ms before the onset of movement, whereas they report being aware of acting around 86 ms before the onset. Thus, these seem to be two clearly distinguishable events from the subjects' point of view.[4]

If experiences of authorship are not conscious intentions, then what else might they be, according to the moderate skeptic? Another temptation is to view them as types of sensory experience—those that accompany one's bodily movements as one engages in action (see Bayne 2011). The main issue with this proposal is that, as I've argued elsewhere (see Mylopoulos 2015, 2017), one and the same bodily movement type can result in an experience of authorship in some cases and not in others. Consider again the movements of the anarchic hand. The activities it engages in, sometimes even involving the manipulation of objects, are often indistinguishable from the third person from intentional actions of the agent. Presumably they are also accompanied by the same types of sensory experience. And yet no experience of authorship accompanies them.

---

[4] Some have argued that participants in these experiments are not directly aware of their decisions to act, but merely infer the timing of their occurrences based on other cues. Banks and Isham (2009) presented participants in a Libet-style task with deceptive feedback in the form of an auditory tone to indicate to participants that they had moved 5 to 60 ms later than they actually had. They found that participants' judgments of when they had moved were also shifted later in time, as a function of the delay in feedback. But the fact that the regular process can be interfered with is not evidence that they do not have direct awareness of decision or intention. And additional work suggests that participants are able to reliably distinguish between awareness of an intention and no awareness of intention in a way that is corroborated by measurements of neural activity (Haggard et al. 1999; see also Schurger et al. 2012, 2016).

In addition, recall that the experience of authorship sometimes arises before bodily movement takes place. When participants in Libet et al.'s (1983) classic study were asked to report the time at which they became aware of performing a spontaneous action, they reported acting on average 86 ms prior to any muscle activity as measured by EMG. This finding fits well with the point made earlier about experiences of authorship involving awareness "from the inside" of oneself as performing an action. Typically, one does not need to observe one's body in order to be cognizant of what action one is performing. In situations where one does, one feels disconnected from oneself as an agent, as in the case of AHS. So we should be looking for a non-sensory experience that takes place prior to the onset of action if we are to subsume the experience of authorship under a more familiar experience type. But we have already ruled out experiences of trying linked to conscious intentions. What to do?

The attempt of the moderate skeptic to reduce the experience of authorship to the very mental states and processes involved in action production, namely intentions and sensory states, seems to have hit a dead end. Why is that? There are two diagnoses available. The first is provided by Shepherd (2017), who maintains that experiences of acting "cannot be described as the co-conscious conjunction of an experience of trying and easily associated perceptual experience of things happening" (p. 432), as the moderate skeptic would like. This is because, as Shepherd argues, the temporal and spatial content of experiences of acting are the result of the robust *integration* of the information carried by both cognitive states (i.e., the agent's intention) and perceptual states. An experience of oneself as performing an action thus has a "distinct unity" that goes beyond the mere co-occurrence of the types of experience that the moderate skeptic appeals to in order to account for it.

This is one way in which one might defend a proprietary view of the experience of authorship. On this view, it cannot be subsumed under familiar experience types, but must be understood as a unified experience of its own, involving rich integration between the experience of trying and the perceptual experience of one's body moving. It is less clear, however, what to say about the experience of controlling an action to varying degrees on this account, since this would seem to require some monitoring or assessment of how well what one's body is doing matches one's intention, and this goes beyond what is on offer here. This consideration might lead one to a second diagnosis of the moderate skeptic's mistake.

The second diagnosis is this: the moderate skeptic's reductive account of agentive experience does not work because the experiences of authoring and controlling an action must involve awareness that carries information *about* the states and processes involved in action initiation and control—and this cannot be provided by those states and processes themselves. This suggests a natural role for *metarepresentation* in both types of experience. In the next and final section, I present a sketch of this approach as a second way of fleshing out the proprietary nature of agentive experiences.

6 Metarepresentation, Metacognition, and Agentive Experiences

For some guidance on how to understand agentive experience as a function of metarepresentation, we can start by looking to higher-order theories of consciousness, which hold that what it is for some mental state to be conscious is for one to be suitably aware of oneself as being in that state. Further, on such views, what it is to be aware of oneself as being in some state is a matter of representing oneself as being in it, either via a higher-order thought (see, e.g., Rosenthal 1986, 2005), or a quasi-perceptual state (see, e.g., Lycan 1996). I will remain neutral here on these two ways of understanding the relevant higher-order state and talk simply in terms of higher-order awareness (HOA).

If one takes on board a higher-order approach to consciousness, one can then say that what it is to have an experience of authorship is to have a HOA of oneself as performing an action, construed as a form of metarepresentation (see Mylopoulos 2017 for exploration of a view on which the source of the relevant HOAs is thoughts). Similarly, one can say that what it is to have an experience of controlling some action is to have some HOA of oneself as doing so, construed along the same lines.

This gives us an attractive strategy for establishing a proprietary phenomenology of agency, since we can hold that just as having a HOA that represents oneself as being in a mental state is sufficient for one to have a subjective experience of being in that mental state (on this point, see Rosenthal 2011, p. 433–434), having a HOA that represents oneself as performing an action is sufficient for generating an experience of performing it. And the experience generated by way of such an HOA would not be reducible to the experience of intending and the experience of one's body moving, although these might accompany it.

I think this view holds promise, and especially so as a way of explaining the experience of authorship in particular. For HOAs of the sort we are considering to always represent

oneself as being in some state or, in this case, being engaged in some process. And so we can easily account for the experience of self as source of an action by way of this self-referential component of the HOA. Once again, though, it is less clear how well this view fares in explaining the experience of control, since here there is an evaluative component involved that monitors the success of one's action execution, and bare states of HOA do not typically play such roles. It is here that I think we can avail ourselves of another sort of metarepresentation, this time the kind that goes under the label of *metacognition*.

As an illustrative example of metacognition, consider what happens when you are asked a question, such as who was the first female winner of the Nobel Prize. In such a case, you will sometimes have a sense that though the answer (Marie Curie) is not available to you in that moment, you will be able to recall it. That sense is a form of metacognition referred to as the 'feeling of knowing' (FOK) (Nelson and Narens 1990). Other common forms of metacognition include judgments of learning (JOL) (Koriat and Levy-Sadot 1999) and the tip-of-the-tongue (TOT) phenomenon (Schwartz 2006).

What is common to all of these forms of metacognition is that a capacity is engaged through which the states or processes of an "online" cognitive system are evaluated or monitored, and there is a subjective feeling attached to this operation. In addition, metacognitive capacities are employed for the purposes of controlling the processes they are about. We can thus follow Shea et al.'s (2014) characterization of metacognition as "control processes that make use of one or more metacognitive representations, that is, representations of a property of a cognitive process" (p. 187).

This characterization not only helps to capture what is common to cases of metacognition in different domains, it helps us distinguish metacognition from HOA more generally (see Rosenthal 2012). Though, on the view we are considering, metacognition, like HOA, is a type of metarepresentation (see also Proust 2013 on procedural metacognition) and it differs from the HOAs posited by higher-order theories insofar as the latter type of awareness need not involve any evaluation of first-order processes, and need not be in the service of any control functions. In many cases of higher-order awareness we're simply aware of the mental states or processes that we are engaged in, such as when one is aware of the mental imagery involved in a daydream, without this involving any further input to a control process. Metacognition, by contrast, is for control.

The suggestion here is that, just as metacognition is often directed toward our perceptual detection processes as well as our mnemonic retrieval process, it can also be directed toward our action execution processes. And when it is so directed, it results in a certain characteristic experience: the experience of control. This proposal offers us yet another way of viewing agentive experience as proprietary, since metacognitive feelings are thought to have their own distinctive phenomenology not reducible to whatever experiences accompany the first-order states and processes they monitor and evaluate.

The experience of control is in fact quite naturally construed as an experience constituted by a metacognitive state that monitors and evaluates one's action performance. Indeed, there have been a number of empirical defenses of this type of view (see, e.g., Metcalfe and Greene 2007; Chambon et al. 2014; Wenke et al. 2010; Carruthers 2015). A study by Wenke et al. (2010) is particularly telling here. In the study, participants were asked to press left or right keys in response to target left or right pointing arrows. Prior to the presentation of the target, they were presented with subliminal arrow primes that were either compatible or incompatible with the target. After participants selected their response (left or right key), with variable delay, they were presented with color patches, the color of which depended on whether the primes were compatible or incompatible with the target. At the end of each block, participants were then asked to rank-order how much control they experienced for each color that appeared. Unsurprisingly, participants' reaction times were faster for compatible primes versus incompatible primes. But most importantly for our purposes, participants also gave higher control ratings for colors following action-compatible primes versus colors following action-incompatible primes.[5] The authors themselves take this to indicate a form of metacognition that monitors the *fluency* of the action selection process. But it can just as well be taken as metacognitive monitoring of the extent to which the action one produces successfully satisfies one's intention: an experience of control.

If the foregoing is correct, then we have a promising way of understanding the proprietary nature of agentive experience. It is not reducible to the experiences

---

[5] In a subsequent study, Chambon et al. (2012) found that participants reported feeling more in control of colors corresponding to prime-compatible actions versus those corresponding to prime-incompatible actions even when reaction times were slower for the former, thus dissociating fluency of action selection from performance monitoring.

accompanying the first-order states and processes that generate action, because it is about those processes. We can think of the experience as the result of different metarepresentational capacities that are typically engaged when an action is being performed: HOA and metacognition. This makes available alternative explanations for the subjective contrast between acting and being passively moved, which fare better than those of both the strong and moderate skeptic.

7. Conclusion

In this paper, I have made a case for the proprietary nature of agentive experience. I have done so by identifying and then confronting two skeptical stances toward such a view, the first on the grounds that experiences of passivity are sufficient to account for the subjective contrast between action and passive movement, and the second on the grounds that agentive experience is reducible to familiar cognitive and sensory phenomenology accompanying action. I have argued that once one appreciates the different facets of agentive experience—in particular the experience of authorship and control—and their metarepresentational nature, these are better construed as a form of HOA and metacognition, respectively, and thus proprietary in just the way the skeptics deny.

References

Banks, W.P. and Isham, E.A. (2009). We Infer Rather than Perceive the Moment We Decided to Act. *Psychological Science*, 20, 17–21.

Bayne, T. (2008). The Phenomenology of Agency. *Philosophy Compass*, 3(1), 182–202.

Bayne, T. (2011). 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.

Bermúdez, J. (2010). Action and Awareness of Agency: Comments on Christopher Frith. *Pragmatics & Cognition*, 18, 584–596.

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.

Carruthers, G. (2015). A Metacognitive Model of the Feeling of Agency Over Bodily Actions. *Psychology of Consciousness: Theory, Research, and Practice*, 2(3), 210–221.

Chambon, V., Filevich, E., and Haggard, P. (2014). What is the Human Sense of Agency, and Is It Metacognitive? In S.M. Fleming and C.D. Firth (eds), *The Cognitive Neuroscience of Metacognition* (pp. 321–342). Berlin: Springer.

Chambon, V., Wenke, D., Fleming, S.M., Prinz, W., and Haggard, P. (2012). An Online Neural Substrate for Sense of Agency. *Cerebral Cortex*, 23(5), 1031–1037.

Csikszentmihalyi, M. (2000). *Beyond Boredom and Anxiety*. San Francisco: Jossey-Bass Publishers.

Danto, A. and Morgenbesser, S. (1963). What We Can Do. *Journal of Philosophy*, 60(15), 435–445.

Davidson, D. (2001). Agency. *Essays on Actions and Events* (pp. 43–61). Oxford: Oxford University Press.

Della Sala, S. (2005). The Anarchic Hand. *The Psychologist*, 18(10), 606–609.

De Vignemont, F. (2007). Habeus Corpus: The Sense of Ownership of One's Body. *Mind and Language*, 22(4), 427–449.

Dow, J.M. (2017). Just Doing What I Do: On the Awareness of Fluent Agency. *Phenomenology and the Cognitive Sciences*, 16(1), 155–177.

Dreyfus, H. (2007). Response to McDowell. *Inquiry*, 50(4), 371–377.

Fried, I., Haggard, P., He, B.J., and Schurger, A. (2017). Volition and Action in the Human Brain: Processes, Pathologies, and Reasons. *Journal of Neuroscience*, 37(45), 10842–10847.

Frith, C.D., Blakemore, S., and Wolpert, D.M. (2000). Explaining the Symptoms of Schizophrenia: Abnormalities in the Awareness of Action. *Brain Research Reviews, 31*(2–3), 357–363.

Gandevia, S.C., Killian, K., McKenzie, D.K., Crawford, M., Allen, G.M., Gorman, R.B., and Hales, J.P. (1993). Respiratory Sensations, Cardiovascular Control, Kinaesthesia and Transcranial Stimulation During Paralysis in Humans. *Journal of Physiology*, 470: 85–107.

Grünbaum, T. (2015). The Feeling of Agency Hypothesis: A Critique. *Synthese*, 192, 3313–3337. https://doi.org/10.1007/s11229-015-0704-6

Haggard, P. and Chambon, V. (2012). Sense of Agency. *Current Biology*, 22(10), R390–R392.

Haggard, P. and Eimer, P. (1999). On the Relation Between Brain Potentials and the Awareness of Voluntary Movements. *Experimental Brain Research*, 126(1), 128–133.

Haggard, P., Newman, C., and Magno, E. (1999). On the Perceived Time of Voluntary Actions. *British Journal of Psychology*, 90(2), 291–303.

Horgan, T. (2007). Agentive Phenomenal Intentionality and the Limits of Introspection. *Psyche*, 13(1), 1–29.

Horgan, T. (2012). From Agentive Phenomenology to Cognitive Phenomenology: A Guide for the Perplexed. In T. Bayne and M. Montague (eds.), *Cognitive Phenomenology* (pp. 57–78). New York: Oxford University Press.

Horgan, T., Tienson, J., and Graham, G. (2003). The Phenomenology of First-Person Agency. In S. Walter and H.-D. Heckmann (eds.), *Physicalism and Mental Causation* (pp. 323–340). Exeter: Imprint Academic.

Koriat, A. and Levy-Sadot, R. (1999). Processes Underlying Metacognitive Judgments: Information-Based and Experience-Based Monitoring of One's Own Knowledge. In S. Chaiken and Y. Trope (eds.), *Dual Process Theories in Social Psychology* (pp. 483–502). New York: Guilford Press.

Kriegel, U. (2015). *The Varieties of Consciousness*. New York: Oxford University Press.

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, 623–642.

Lycan, W. (1996). *Consciousness and Experience*. Cambridge, MA: MIT Press/ Bradford Books.

Marcel, A. (2003). The Sense of Agency: Awareness and Ownership of Action. In J. Roessler and 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.

Mellors, C.S. (1970). First-Rank Symptoms of Schizophrenia. *British Journal of Psychiatry*, 117, 15–23.

Metcalfe, J. and Greene, M.J. (2007). Metacognition of Agency. *Journal of Experimental Psychology: General*, 136(2), 184–199.

Mylopoulos, M. (2015). Agentive Awareness Is Not Sensory Awareness. *Philosophical Studies*, 172(3), 761–780.

Mylopoulos, M. (2017). A Cognitive Account of Agentive Awareness. *Mind & Language*, 32(5), 545–563.

Mylopoulos, M. and Shepherd, J. (2020). The Experience of Agency. In U. Kriegel (ed.), *The Oxford Handbook of the Philosophy of Consciousness* (pp.164–187). Oxford: Oxford University Press.

Nelson, T.O. and Narens, L. (1990). Metamemory: A Theoretical Framework and New Findings. In G.H. Bower (ed.), *The Psychology of Learning and Motivation* (pp. 125–173). New York: Academic Press.

Pacherie, E. (2008). The Phenomenology of Action: A Conceptual Framework. *Cognition*, 107, 179–217.

Prinz, J.J. (2012). *The Conscious Brain: How Attention Engenders Experience*. Oxford: Oxford University Press, Kindle Edition.

Proust, J. (2013). *The Philosophy of Metacognition: Mental Agency and Self-Awareness*. Oxford: Oxford University Press.

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.M. (2011). Exaggerated Reports: Reply to Block. *Analysis*, 71(3), 431–437.

Rosenthal, D.M. (2012). Higher-Order Awareness, Misrepresentation and Function. *Philosophical Transactions of the Royal Society: London B: Biological Sciences,* 367(1594), 1424–1438.

Schlosser, M.E. (2011) Agency, Ownership, and the Standard Theory. In J.H. Aguilar, A.A. Buckareff, and K. Frankish (eds.) *New Waves in Philosophy of Action* (pp. 13–31). London: Palgrave Macmillan.

Schurger, A., Mylopoulos, M., and Rosenthal, D. (2016). 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*, 109(42), E2904–E2913.

Schwartz, B.L. (2006). Tip-of-the-Tongue States as Metacognition. *Metacognition Learning*, 1, 149–158.

Shea, N., Boldt, A., Bang, D., Yeung, N., Heyes, C., and Frith, C.D. (2014). Supra-Personal Cognitive Control and Metacognition. *Trends in Cognitive Sciences*, 18(4), 186–193. https://doi.org/10.1016/j.tics.2014.01.006

Shepherd, J. (2016). Conscious Action/Zombie Action. *Noûs*, 50(2), 419–444.

Shepherd, J. (2017). The Experience of Acting and the Structure of Consciousness. *Journal of Philosophy*, 114(8), 422–448.

Spence, S.A., Brooks, D.J., Hirsch, S.R., Liddle, P.F., Meehan, J., and Grasby, P.M. (1997). A PET Study of Voluntary Movement in Schizophrenic Patients Experiencing Passivity Phenomena (Delusions of Alien Control). *Brain*, 120(11), 1997–2011.

Strawson, G. (2010). *Freedom and Belief* (rev. edn.). Oxford: Oxford University Press.

Velleman, J.D. (1992). What Happens When Someone Acts? *Mind*, 101, 461–481.

Wenke, D., Fleming, S.M., and Haggard, P. (2010). Subliminal Priming of Actions Influences Sense of Control over Effects of Action. *Cognition*, 115(1), 26–38.</REF>