

Agentive awareness is not sensory awareness

Myrto I. Mylopoulos


Philosophical Studies

An International Journal for Philosophy
in the Analytic Tradition

ISSN 0031-8116

Philos Stud

DOI 10.1007/s11098-014-0332-x



Your article is protected by copyright and all rights are held exclusively by Springer Science +Business Media Dordrecht. This e-offprint is for personal use only and shall not be self-archived in electronic repositories. If you wish to self-archive your article, please use the accepted manuscript version for posting on your own website. You may further deposit the accepted manuscript version in any repository, provided it is only made publicly available 12 months after official publication or later and provided acknowledgement is given to the original source of publication and a link is inserted to the published article on Springer's website. The link must be accompanied by the following text: "The final publication is available at link.springer.com".

Agentive awareness is not sensory awareness

Myrto I. Mylopoulos

© Springer Science+Business Media Dordrecht 2014

Abstract In this paper, I argue that the conscious awareness one has of oneself as acting, i.e., agentive awareness, is not a type of sensory awareness. After providing some set up in Sect. 1, I move on in Sect. 2 to sketch a profile of sensory agentive experiences (SAEs) as representational states with sensory qualities by which we come to be aware of ourselves as performing actions. In Sect. 3, I critique two leading arguments in favor of positing such sensory experiences: the argument from pathology and the argument from cognitive impenetrability. Since neither of these arguments succeeds, the case for positing SAEs is dealt a significant blow. I proceed in Sect. 4 to advance my positive argument against SAEs. The argument runs as follows: If SAEs exist, then they must exist in some sensory modality or set of sensory modalities. Either the relevant sensory modalities are ones that we already recognize, or they are novel sensory modalities. I will argue that neither of these options is workable, and so we have nowhere to locate SAEs. Agentive awareness is not sensory awareness.

Keywords Sense of agency · Agentive awareness · Phenomenology · Consciousness · Comparator model · Action theory

1 Introduction

A central feature of our mental lives is the conscious awareness we have of ourselves as acting. Following Bayne and Pacherie (2007), I will call this *agentive awareness*.¹ One of the central questions surrounding agentive awareness is the

¹ This brand of self-awareness has gone under a number of other labels as well, most popularly, “the sense of agency” (e.g., de Vignemont and Fourneret 2004; Gallagher 2000; Marcel 2003; Peacocke

M. I. Mylopoulos (✉)
Philosophy Department, Fordham University, 527 E. 6th St., Apt. BW, New York, NY 10009, USA
e-mail: myrto.mylopoulos@gmail.com

question of what *kind* of awareness it is. In general, one is made aware of things by being in appropriate mental states. One's perceptual states, for example, make one aware of one's body and physical environment. And similarly, one is sometimes aware of something by having a thought about it as being present. But where in our mental architecture—thought, perception, or elsewhere—does *agentive* awareness lie? That question will be the focus of this paper.

It is plain that one sometimes has thoughts or makes judgments to the effect that one is acting. But rarely do theorists defend the view that agentive awareness is, at bottom, a matter of having suitable judgments about one's own occurrent agency. Rather, it is commonly held that such judgments are based on some *antecedent* awareness of oneself as acting. And the view most frequently endorsed is that such antecedent awareness is best understood as a form of *experience*. For example, Bayne and Pacherie (2007) write:

We hold that the 'vehicles' of agentive self-awareness are often more primitive than judgments. Think of what it is like to push a door open. One might judge that one is the agent of this action, but this judgment is not the only way in which one's own agency is manifested to oneself; indeed, it is arguably not even the primary way in which one's own agency is manifested to oneself. Instead, one *experiences* oneself as the agent of this action (476).

But this just pushes the question back a step, for now we must ask what kind of *experiences* constitute agentive awareness? After all, the term 'experience' is often used simply to describe a mental state with some phenomenal character, and this could apply to a wide variety of mental state types, including sensations, perceptions, emotions, and even thoughts (cf. Pitt 2004). In what follows, I will focus on the natural proposal that agentive experiences are best understood as forms of *sensory* experience and I will argue that this proposal does not succeed. While the position I defend in this paper is thereby strictly a negative one, if I am correct, then we have some clear motivation for the positive project of developing non-sensory accounts of agentive awareness.

A version of the sensory approach has recently been defended by Bayne (2011), according to whom:

Just as we have sensory systems that function to inform us about the distribution of objects in our immediate environment, damage to our limbs, and our need for food, so, too, we have a sensory system (or systems), whose function it is to inform us about facets of our own agency (356).

On Bayne's view, such experiences arise out of a dedicated sensory modality. But a proponent of the view that agentive experiences are sensory experiences need not follow Bayne in claiming this. In addition to Bayne's view, I will consider another popular version of the sensory approach, as defended by Frith (1992, 2007) and his colleagues, that does not take this commitment on board.

Footnote 1 continued

2003). Others include "the phenomenology of agency" (e.g., Pacherie 2008), "control consciousness" (e.g., Mandik 2010), and "action consciousness" (e.g., Prinz 2007).

The paper is organized as follows: In Sect. 2, I lay out the characteristics that agentive experiences must have if they are to be understood as sensory experiences. In Sect. 3, I critique two of the main arguments that are offered in favor of positing this class of sensory experiences. In Sect. 4, I advance my positive argument against the view that agentive experiences are sensory experiences. The argument runs as follows: If agentive experiences are sensory experiences, then they must exist in some sensory modality or set of sensory modalities. This sensory modality or set of sensory modalities must either be one that we already recognize (e.g., vision, hearing, touch, smell, taste, and proprioception) or novel. I argue that neither option is workable by way of showing that leading accounts for each option fail. In the absence of any modality within which to plausibly locate sensory agentive experiences (SAEs), I conclude that there are none.

2 What are sensory agentive experiences (SAEs)?

To be clear, in denying that there are SAEs, I am not denying, naturally, that we sense our actions. We regularly sense our bodily movements, and these are often themselves actions. Nor am I denying that engaging in action sometimes *shapes* aspects of our sensory experience. There is some evidence, for example, that events following our actions appear subjectively to occur closer in time than those following our non-actions, even if the temporal interval is, in fact, the same—a phenomenon known as the “intentional binding effect” (see Haggard et al. 2002). I am also not contesting that there is a proprietary *phenomenology* of agency. Even if there are no SAEs, this leaves open the possibility that acting has a distinctive phenomenology that is tied to non-sensory states. What I do dispute, however, is that we come to be consciously aware of ourselves *as* acting by way of any sensory experience or set of sensory experiences with the features I will describe in this section.

Broadly speaking, agentive experiences, whether sensory or not, are thought to have two main features. First, they are viewed as representational states (e.g., Bayne 2011; Bayne and Pacherie 2007; Carruthers 2012; Pacherie 2008; Synofzik et al. 2008). They are not construed as purely “raw feels” without representational properties. In this respect, they are on par with other representational states that we readily recognize, such as perceptions, beliefs, and desires. Second, as I indicated earlier, agentive experiences are often supposed to be the basis for agentive judgments. When one judges that one is acting, it is commonly thought that one does so *because* one is having an agentive experience (e.g., Bayne and Pacherie 2007; Bermúdez 2010; Synofzik et al. 2008). Given this putative role in grounding agentive judgments, agentive experiences are most reasonably construed as the way by which we *come to be* aware of ourselves as acting. They are not *merely* states that provide such awareness—agentive judgments satisfy that condition as well—they are the *first stop* for agentive awareness, the primary way by which we become so aware.

This last point is worth highlighting for the reason that, even if it turns out that there are sensory experiences that represent oneself as performing an action, this is

not enough to establish that agentive awareness is sensory awareness, for these sensory experiences might derive from some antecedent agentive awareness that is non-sensory in form. Similarly, the fact that one sometimes forms judgments about the colors of certain objects does not, of course, entail that color awareness is at bottom a matter of forming suitable judgments, since these judgments are themselves based on antecedent color sensations. It is the states by which we first come to be agentially aware that are of present concern, not states further downstream in our psychological processing that are based on these states.

So far I have identified two general features that theorists typically attribute to agentive experiences. But if agentive experiences are *sensory* experiences in particular, then they must also involve sensory qualities. In an attempt to stay relatively neutral with respect to competing theories on the nature of sensory qualities, I will adopt a fairly minimal view of sensory qualities according to which they have two main features: (i) they are associated with representational states that are the outputs of sensory modalities, and (ii) they correspond to sensory properties of objects in the world (e.g., color, shape, texture), such that when a healthy creature senses these objects under normal conditions, the relevant sensory modality outputs a state with the relevant sensory quality or qualities in virtue of which the object appears as it does. So when one has a conscious visual experience of, say, a ripened lime, one's experience involves green sensory qualities in virtue of which the lime appears as it does. Similarly, when one has a conscious agentive experience, if these experiences are sensory states, they must involve sensory qualities in virtue of which one's acting appears as it does. Indeed, such sensory qualities are arguably among the features that distinguish SAEs, if they exist, from the agentive judgments that they are supposed to support.

Some hold that the sensory qualities associated with experiences are to be identified with the representational properties of those experiences (e.g., Harman 1990; Tye 2000). Others maintain that sensory qualities are something over and above the representational properties of experience (e.g., Block 2003; Chalmers 2004). I will not take a stance on this issue here. Either way, if agentive awareness is sensory awareness, we should be able to identify some type of sensory quality or set of sensory qualities uniquely associated with it. A type of sensory quality that sometimes accompanies awareness of oneself as acting and at other times accompanies awareness of oneself as passively moving will not do here.

The goal of this paper is, then, to argue that SAEs, understood as states with these three features, do not exist. There is no type of mental state that represents oneself as performing an action, possesses sensory qualities, and is the primary way by which we come to be aware of ourselves as acting. Of course, not all are convinced that this is the case. I turn now to consider two arguments that support the positing of SAEs.

3 The arguments from pathology and cognitive impenetrability

It might strike one as odd to require *arguments* to establish the existence of a certain kind of experience. After all, one might reason, all experiences are directly introspectible, and if so, what good are arguments? As Searle (1983) remarks:

It is a bit difficult to know how one would argue for the existence of perceptual experiences to someone who denied their existence. It would be a bit like arguing for the existence of pains: if their existence is not obvious already, no philosophical argument could convince one (44).

When it comes to agentive experiences, however, many have supposed that they are *not* readily available to introspection, at least to the same degree as familiar sensory experiences like pains and smells. They are, instead, frequently described as “thin” and “evasive” (Metzinger 2006).

Some may embrace a skeptical conclusion on this basis, holding that the elusive character of agentive experiences should lead us to deny their existence. But this is too quick. Even if we were to lack *complete* introspective access to agentive experiences, this would not entail that they do not exist. This would follow only if introspection gives us exhaustive access to our mental lives, and we have good independent reason to reject this Cartesian conviction, on the basis of empirical findings such as subliminal priming and blindsight. Moreover, we have other ways of becoming aware of our experiences that do not rely on directly introspecting them. For example, we can establish their existence inferentially (cf. Tye 2003). I might infer from my awareness of myself as pouring coffee that I am having an SAE to that effect. And I could do this even if I could not directly introspect any features of that experience. So introspection neither confirms nor disconfirms the existence of SAEs. We must look instead to theory. In the remainder of this section, I thus consider two arguments—the first in favor of positing agentive experiences more generally, and the second in favor of positing SAEs in particular.

3.1 The argument from pathologies of agency

A popular line maintains that sound explanations for certain pathologies of agency make a significant appeal to agentive experiences, and that this gives us good reason to posit them.² Consider anarchic hand syndrome, which involves lesions to the motor cortex, and drives individuals to perform what *seem*, from a third-person point of view, to be actions with one of their limbs, all the while denying authorship of its movements. It will be helpful to illustrate the nature of this condition with an episode that the neurologist Sergio Della Sala (2005) describes concerning an individual with anarchic hand syndrome with whom he works:

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.

² Though this first argument, if it were successful, would not by itself be sufficient for establishing the existence of *sensory* agentive experiences, it is worth considering here, since some theorists may be tempted to appeal to it in order to move towards that conclusion.

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' (606).

As this episode makes clear, Mrs. GP does not identify with the movements of her anarchic limb, no matter how much they appear to be her actions from the outside. She is aware of them as being out of her control, and does not take them to express her own agency.

Regarding such cases, Bayne (2011) writes:

It seems plausible to appeal to agentic 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 (360, emphasis in original).

Bayne is urging here that a reasonable explanation for why anarchic hand patients like Mrs. GP deny ownership of their anarchic movements appeals to experiences of passivity replacing regular agentic experiences. This diagnosis is problematic, however, at least if it is meant to support the claim that agentic experiences themselves exist. First, as Bayne himself recognizes, the existence of experiences of passivity does not give us strong reason to suppose that there are agentic experiences, nor agentic awareness of any form for that matter. This would be like inferring on the basis of experiences of pain that whenever one is not having such experiences, one is having experiences of painlessness. Instead, it may be that whenever something goes wrong with our actions, we come to have experiences of passivity, but that when things are going smoothly, no corresponding experience of acting is present.

Second, even if the existence of experiences of passivity were to give us reason to believe in the existence of agentic experiences, it would not follow that either of these experiences is sensory in form. It may be that neither of these experiences is sensory. Or it may be that experiences of passivity are sensory while agentic experiences are not.

On this last point, we may distinguish between two possible relationships between agentic experiences and experiences of passivity (cf. de Vignemont's (2011) discussion of the relationship between the sense of ownership and the sense of disownership). On the first possibility, experiences of passivity derive from psychological mechanisms that are independent of those responsible for agentic experiences. If this were the case, then facts about experiences of passivity would not entail facts about the nature of agentic experiences. It may be that the psychological mechanism responsible for the former is sensory in nature, and that the mechanism responsible for the latter is not. And until we have independently determined that this first possibility does not, in fact, hold, we are not licensed to make any inferences from claims about one to claims about the other.

On a second way of modeling the relationship between experiences of passivity and agentic experiences, they are generated by the same psychological mechanism, such that, depending on the outcome of the relevant operations of that mechanism, it

may yield either type of experience. This possibility makes it more likely that, if experiences of passivity are sensory, then so too are agentive experiences, since one and the same psychological mechanism is more likely than not to produce outputs of one and the same mental type. But again, until we have independently determined that these two types of experience *do* share the same psychological mechanism, we are not licensed to draw any inferences from claims about one to claims about the other.

A further difficulty with the argument from pathologies of agency is the claim that denials of action authorship in anarchic hand syndrome are reasonably explained by supposing that experiences of passivity have replaced agentive experiences. Such a move is only useful for establishing agentive experiences if it is the only or best explanation we have available to us. But this is not the case, since a number of equally good alternative explanations are available for what underlies denials of action authorship in anarchic hand syndrome, which make no appeal to the absence of agentive experiences, nor to the presence of passivity experiences. For example, these individuals are typically aware that they are unable to inhibit the movements of their anarchic limb, as well as aware that what their limb does clashes with what they consciously intend or want it to do. On the basis of these observations, it is not difficult to understand why they deny that the actions of the limb are their own. Neither passivity experiences, nor the absence of agentive experiences need to play a role in this equally credible alternative explanation of the psychology underlying the reports of individuals with anarchic hand. And if so, we are left without a good reason to posit them here.

3.2 The argument from cognitive impenetrability

Another type of argument sometimes proffered for the existence of SAEs appeals to features of agentive awareness that are allegedly characteristic of perception. For instance, Bayne (2011) argues that agentive awareness is sensory awareness on the grounds that such awareness is cognitively impenetrable, and cognitive impenetrability is a hallmark of perception. He writes: “Perhaps judgment can penetrate agentive experience under some conditions, but it seems clear that on the whole agentive experience exhibits the kind of doxastic impenetrability that is characteristic of perception” (361).

As evidence for the cognitive impenetrability of agentive experience, Bayne (2011) focuses again on the case of anarchic hand syndrome. He reasons that if an anarchic hand patient were to come to believe that the behavior of her anarchic limb were, in fact, her own, this would not suffice to recover her agentive experiences with respect to this behavior. Bayne urges:

The patient can no more restore the missing experience of agency by forming the belief that her anarchic actions are truly her own than you or I can correct our visual experiences of the Müller-Lyer illusion by forming the belief that the lines in question are equal in length (360).

While this may be a tempting argument, there are reasons not to accept it. First, we lack independent empirical support for the claim that an anarchic hand patient would continue to lack agentive awareness after coming to sincerely believe that she

is the agent of her anarchic actions. It may be instead that were she to form the relevant belief in a suitable way, her agentic awareness would indeed be restored. We need some further evidence or theoretical considerations to settle the issue.

Moreover, it is not obvious that if some form of awareness is cognitively impenetrable, it must thereby be sensory awareness.³ Suppose, for example, that one were given a lengthy mathematical proof by someone viewed as an authority in this domain for the claim that one plus one does not, in fact, add up to two. And suppose further that one accepted this proof, thereby adopting the belief that it is not the case that one plus one is equal to two. What would then happen to one's conflicting belief, acquired early on in childhood and frequently reinforced over the years, that one plus ones adds up to two? It is not unreasonable to suppose that given the strength of this belief, it would continue to persist. In particular, it might continue to *seem* that one plus one is two despite one's also accepting the mathematical proof that this is not the case, and holding a contradictory belief to that effect. And if holding some belief does not eradicate its inconsistent counterparts—in other words, if some inconsistent beliefs are cognitively impenetrable relative to each other—it does not, of course, follow that they are sensory states.

Applying this principle to the anarchic hand case, it might be that the individual comes to believe, perhaps also on the basis of authoritative testimony, that she is the agent of the anarchic limb's behavior, but that this belief gets “held in check” by a stronger belief that she is not the agent, such that it does not sway her overall assessment of the situation (cf. Armstrong 1968; Byrne 2009). Having the deliverances of some form of awareness persist despite the presence of inconsistent beliefs does not give us good reason to conclude that the awareness in question is sensory in nature, as the argument from cognitive impenetrability insists.

In this section I have challenged two leading arguments that make up the case for positing SAEs. In the absence of further arguments in support of their existence, and the lack of any clear introspective access to such experiences, the case for positing them becomes increasingly thin. Of course, none of this yet shows that there are no SAEs, just that we do not, at present, have good reason to think so. I move on now to advancing positive considerations against the existence of SAEs.

4 The case against SAEs

All sensory experiences belong to some sensory modality. When it comes to SAEs, then, there must be some associated modality as well. There are two different strategies for identifying the relevant modality: reductive and non-reductive. On a reductive strategy, agentic experiences are identified with sensory experiences in one or more of the familiar sensory modalities. Non-reductive strategies, on the

³ Certainly there is considerable dispute that the converse holds. For example, MacPherson (2012) argues that there is at least one case of color perception in which one's beliefs about the typical colors of objects affects the colors that those objects appear to have, and that the interpretation of this result that appeals to cognitive penetration of color perception cannot be dismissed.

other hand, treat agentive experiences as *sui generis* experiences that exist within a novel sensory modality dedicated to sensing our own agency.

In this section, I will consider the leading versions of each type of strategy. While taking different approaches, the accounts I evaluate here share in common the core idea that SAEs are intimately connected to the processes underlying fine-grained sensorimotor control. Indeed, they both adopt a particular theoretical model of how such control operates known as the comparator model (Blakemore and Frith 2003; Frith 1992; Wolpert 1997). They then seek to account for SAEs in terms of specific components of this model.

For present purposes, the important component of the comparator model is the so-called *forward model*. Before an agent performs an action, the forward model is thought to take as input (i) a copy of the motor command generated by the motor system, which specifies the relevant details of the bodily movement to be performed in the service of the agent's goal, and (ii) the estimated current state of the central nervous system. On the basis of these two representations, it outputs predictions of the sensory consequences of the motor command, both with respect to the sensory properties of the bodily movement itself (e.g., moving one's finger), and the sensory properties of the effects of the bodily movement (e.g., the lights going off). These predictions are then compared with the agent's goal, as well as sensory feedback from the action and its consequences to yield a match or a mismatch (Blakemore and Frith 2003; Pacherie 2008). This information subsequently aids in the overall control of the action. As we shall see, on both accounts I will now consider, it is this latter comparison process—between the forward model prediction and sensory feedback from the movement and its consequences—that is the purported locus of SAEs.

4.1 Against a reductive account of SAEs

If SAEs are identified with experiences in the familiar sensory modalities, then the best candidate modality is proprioception—the sense we have of our limb positions and movements.⁴ After all, it is through vision and proprioception that we primarily sense our actions, but we do not always see what our bodies are doing, so that leaves proprioception as the dominant action-sensitive modality. Perhaps, then, SAEs are to be identified with proprioceptive experiences.

While tempting, there is a significant problem with this proposal. In some cases, agentive awareness diverges with respect to bodily movements that are the same. In other words, one may be aware of the same type of bodily movement as one's own action in one instance, and not in another. Consider, for example, the findings of the neurosurgeon Wilder Penfield, who in the 1950s electrically stimulated the motor cortex of his patients during surgery. This led these individuals to make bodily

⁴ I set aside for now mental actions such as mentally multiplying twelve times thirteen or imagining a pink elephant. It is worth noting, however, that sensory approaches to agentive awareness have difficulty accounting for such actions, especially when grounded in the comparator model, since it is unclear that forward modeling takes place in these cases given that there is no reason for a motor command to be generated (though see Campbell 1999).

movements that were “smooth... involving coordinated sequences of the operation of multiple muscles, which looked to have the character of voluntary actions, at least from the outside” (as reported by Wegner 2002). And yet, the patients denied authorship of their movements, making remarks like: “I didn’t do that. You did” (Penfield 1975).

For an account on which agentic experience is a type of proprioceptive experience, these results are difficult to accommodate. Because the bodily movements that Penfield’s patients disowned were purportedly indistinguishable from the third-person from actions, there is reason to think that they were accompanied by proprioceptive experiences no different from those that accompanied bodily movements for which they presumably *did* have agentic awareness. But if so, then there is no way to explain, on this kind of account, why they would lack agentic awareness in the one case and not in the other. This thereby casts significant doubt on the claim that agentic experiences are properly construed as proprioceptive experiences.

One might object to this line of reasoning on the grounds that bodily movements of the same type need not be accompanied by corresponding proprioceptive experiences of the same type. Consider, by analogy, well-known color context effects in visual perception (see Albers 2006). Here, one and the same stimulus might appear to be different shades of color depending on adjacent stimuli or the background against which it is presented. So the same color stimulus may result in different color experiences. Perhaps proprioceptive experiences corresponding to bodily movements are similarly modulated depending on the context in which a particular bodily movement unfolds. If so, then one should expect to sometimes find that, given two bodily movements A and B of the same type, the proprioceptive experiences corresponding to A are agentic, while those corresponding to B are not.

Indeed, a proposal by Frith might help fill in the picture further here. In a series of experiments, Frith and his colleagues have shown that the sensory effects of our actions are attenuated relative to the effects of our non-actions (Blakemore et al. 1999; Blakemore et al. 1998; Blakemore et al. 2000). Attenuation may be understood as a comparatively lower degree of subjective intensity along one or more of a sensory quality’s dimensions, e.g., the loudness of an auditory sensory quality. In one such experiment, a first condition (“Self-generated”), had participants apply a touch to the palm of their hand using a metal rod with a foam attachment. In a second condition (“Externally Generated”), the experimenter applied the touch to the same spot using the rod. The participants were asked to rate how tickly, intense, and pleasant the touch felt in each condition. Participants consistently rated the self-generated touch to be less tickly, intense, and pleasant than the externally generated touch. In other words, participants judged the effects of their own actions to be attenuated relative to effects that were not the results of their own actions (Blakemore et al. 1999).

The explanation offered for this finding is that since the self-generated touch is the result of the participant’s own action, a forward model prediction is formed on the basis of a motor command and then compared against the sensory consequences of the action, i.e., the tactile sensation (Blakemore et al. 2000). Since there is a match between the two, the sensation is thereby attenuated or subjectively

dampened. In the case of the externally-generated touch, however, no motor command is involved and so no forward model prediction is formed. Therefore, no comparison and no attenuation takes place.

In a separate experiment, schizophrenic individuals with what are known as delusions of alien control were also tested on this task. In delusions of alien control, individuals experience their actions as not being under their control, lacking agentive awareness, and often attributing them to external agents. For instance, one individual suffering from such delusions reported that “[m]y grandfather hypnotized me and now he moves my foot up and down” (Frith et al. 2000, p. 358), while another insisted that “[t]he force moved my lips. I began to speak. The words were made for me” (Mellors 1970, p. 18). Yet another individual, having just made an arm movement, explained: “I felt like an automaton, guided by a female spirit who had entered me during it” (Spence et al. 1997, p. 2001).

Importantly, it has been found that schizophrenic individuals with delusions of alien control do *not* experience sensory attenuation, at least to the same degree as non-schizophrenic individuals and schizophrenic individuals without delusions of alien control. Blakemore et al. (2000) asked diagnosed schizophrenics with and without delusions of alien control, as well as healthy controls, to perform the same self-tickling task. They found that those with delusions of alien control did not rate externally produced tactile stimuli as more “tickly” than those resulting from self-generated movements. By contrast, the schizophrenic individuals without delusions of alien control, and the healthy control participants did rate them as more “tickly”. Other studies have found evidence of a similar lack of sensory attenuation for action effects among schizophrenic individuals (e.g., Lindner et al. 2005; Shergill et al. 2005).

On the basis of these results, some theorists have suggested that sensory attenuation, tied to the workings of the forward model, can be invoked to explain delusions of alien control in schizophrenia. The idea is that *because* sensory experiences of action are not attenuated for schizophrenic individuals, they feel like passive movements for those individuals. In this vein, Frith (2007) writes:

I believe we can now achieve some understanding of [schizophrenic individual] PH's experiences because of what we have discovered about the brain. In our normal state we are hardly aware of the sensations that occur whenever we move. This is because our brain can predict these sensations and suppress our awareness of them. But what would it be like if something went wrong with the prediction and we became aware of the sensations? Normally I am only aware of the sensations when someone else moves my hand. Such a brain abnormality could explain why PH feels as if her arm is being moved by someone else. She is abnormally aware of her bodily sensations when she moves her hand. For her it really does feel as if someone else were moving her hand (109).

At first blush, this proposal seems to help get around the problem raised by the Penfield cases, since it appears to lend support to the claim that the same movements need not be accompanied by the same proprioceptive experiences. If a forward model prediction is formed and a match registered with the sensory feedback from the movement, the accompanying proprioceptive experiences will be attenuated. But if a forward model prediction is not formed, as presumably was the case with

Penfield's patients, a different, non-attenuated sensory experience will result despite type-identical bodily movements. On this view, agentive experiences are themselves types of proprioceptive experience—depending on whether they have been successfully predicted by the forward model and thereby attenuated.

Upon closer examination, however, the issue is not so easily resolved. The problem is that, at most, the results from the self-tickling experiments establish that sensory attenuation applies to the *effects* of our actions, i.e., the touch caused by moving the rod. They do not establish that such attenuation occurs for the bodily actions themselves, i.e., the moving of the rod, and the corresponding proprioceptive experiences. And if so, then we are left without an underlying SAE to explain cases of agentive awareness that accompany purely bodily actions without any salient effects. Since many of the actions for which we experience agentive awareness are of this type, this is a serious shortcoming of the proposal.

However, a separate study by Blakemore et al. (2003) purports to lend support to the claim that bodily actions are themselves accompanied by attenuated proprioceptive experiences. In this study, participants were hypnotized, and underwent PET scans in three different conditions. In the first condition (“Active Movement”), participants were told to move their left arm up and down. In the second condition (“Real Passive Movement”), a pulley system moved their left arm up and down—participants contributed nothing to this movement. In the third condition (“Deluded Passive Movement”), participants were *told* that their arms were being moved up and down passively by the pulley system, but in fact the participants themselves were making the movements. The authors found that in the Deluded Passive Movement condition, there was a significant increase in brain activation in the cerebellum, and in the parietal lobe within the secondary somatosensory cortex and the inferior parietal cortex, compared with the Active Movement condition.⁵ In other words, there was increased activation during movements that were believed to be passive even though they were not, compared with movements that were believed to be actions.

It is tempting to interpret these results as showing that sensory attenuation occurs for the sensory experiences accompanying our bodily actions themselves; the movements that are believed by participants to be actions are attenuated, whereas the movements that are believed to be passive are not. But while the results are suggestive, there are good reasons to resist such an interpretation.

One worry is that it relies on the questionable assumption that greater brain activation in particular regions corresponds to greater intensity of conscious sensation. This assumption is not warranted without a far better understanding than we currently possess of how brain activation relates to the subjective intensity of conscious experience.

Second, the effects of hypnosis on the brain are still poorly understood, and the relationship between hypnosis and conscious experience is even less well understood. Until we have progressed further in our understanding of this phenomenon, it is methodologically dangerous to place much weight on results from hypnosis studies alone.

⁵ I'm grateful to an anonymous reviewer for directing me to this study.

But perhaps most importantly, even if we set these two worries aside, another study conducted by Weiller et al. (1996) found conflicting results. Weiller et al. had participants undergo PET scans in two conditions: an active one in which they flexed and extended their elbows, and a passive one in which their elbows were flexed and extended by a torque motor. The experiments found the same activation in primary sensorimotor cortex during both active and passive movements. The authors conclude: “Activation was almost identical in location and amount in the primary sensorimotor cortex during both passive and active movements” (108).

If what I have been arguing is correct, we do not yet have any solid evidence that sensory attenuation accompanies our bodily actions themselves and not just their effects. But even if we did have such evidence, it is worth noting that this would be consistent with the causal relationship running in the other direction: actions and their effects may feel subjectively attenuated as a *result* of antecedent agentive judgments. And if so, then it is unlikely that subjectively attenuated proprioceptive experiences would be the *basis* for such judgments, as is commonly assumed. Indeed, a recent study by Desantis et al. (2012) would support just this hypothesis. In the study, participants were asked to judge the loudness of an auditory tone that they believed was either self-caused or externally-generated by way of key presses. In actuality, the tones were always triggered by the participants’ key presses. But even so, participants judged the tones to be attenuated when they believed that they were self-caused rather than externally caused, suggesting that this effect is driven by one’s beliefs about the causal relationship between one’s action and its effects. And if so, then even if attenuation were found for actions themselves, agentive judgments would likely be responsible for these experiences rather than the other way around.

One might reply here that the causal relationship could run in both directions: agentive judgments could modulate the attenuation of proprioceptive experiences, and attenuated proprioceptive experiences, in turn, could reinforce agentive judgments. Even if this were the case, though, sensory attenuation would still not be the primary source of agentive awareness, in the sense outlined earlier; it would be something that bolsters or intensifies our antecedent agentive awareness of ourselves as acting, but does not provide the basis for such awareness. In addition, while the Desantis et al. (2012) study provides evidence for agentive judgments modulating sensory attenuation, we do not yet have clear evidence for sensory attenuation affecting agentive judgments.

Another problem with Frith’s proposal is that credible alternative explanations that do not appeal to forward modeling are available in order to account for the lack of sensory attenuation among schizophrenic individuals. For example, schizophrenic individuals have been found to have attentional deficits, both in selective attention and sustained attention (Gold et al. 2007; Carter et al. 2010). Such deficits could impair their ability to properly allocate attentional resources to the tactile stimuli in the self-tickling task, thus corrupting the results. In particular, they may not differentially attend to self-generated stimuli, which are expected and thereby typically attended to less, as compared with externally-generated stimuli, which are less expected and thereby typically attended to more. This would explain why they do not rate the tickliness of externally-generated sensations to be higher than self-

generated ones; they simply allocate roughly equal amounts of attentional resources to each and thereby experience their qualities as being subjectively on par with one another (cf. Mylopoulos 2012).

These difficulties make it unlikely that sensory attenuation can save a reductive account of SAEs in terms of proprioception. But one might be tempted to argue here that, even if proprioception alone cannot support SAEs, proprioception *plus* vision can do so. After all, though it is true that we do not see all of our actions, we do *typically* have visual feedback pertaining to our bodily movements. Perhaps, combined with proprioception, this might yield multimodal SAEs.

Though this proposal may be inviting, it does not hold up. The perceptible properties of our movements accessible by vision are even less informative than proprioception with respect to whether the movements in question are actions or not. There is no visual quality or set of visual qualities that corresponds to actions as against non-actions. But if not, and if, as I have just argued, there is no such quality or set of qualities in proprioception either, then it is doubtful that vision and proprioception will somehow jointly result in SAEs.

Before moving on from the reductive strategy for identifying SAEs, it is worth considering an interesting variant of the approach, defended by Prinz (2007).⁶ On Prinz's view, agentic experiences are forward model predictions, which he construes as sensory images.⁷ His reason for viewing them as sensory images is that he holds that in order for predictions to be compared with reafferent sensory feedback, they must have the same representational format as this feedback, so they must be sensory states as well. And if so, then the most plausible option is that they are a type of sensory imagery, which Prinz understands as the "willful reactivation" of dedicated input systems.⁸

If the comparator model is an accurate model of sensorimotor control, and Prinz is right in viewing forward model predictions as sensory images, then his proposal does identify a type of sensory experience that accompanies our actions but not our non-actions. And in this way it may seem to identify a solid candidate for the elusive SAE. There is a significant problem with Prinz's view, however (for further discussion, see Mylopoulos 2011; Mandik 2010). Forward modeling has the function of anticipating bodily movements resulting from motor commands. In cases where an action is being successfully executed, forward model predictions will match sensory feedback from the bodily movement. But in some cases, in

⁶ Though see also Prinz (2012, pp. 237–239) in which he notes some problems for the Prinz (2007) view, and ends up abandoning it in favor of a deflationary approach to agentic awareness on which there is no experience of authorship, but simply experiences of our actions and their disruption.

⁷ It is worth noting that Prinz's (2007) original proposal is that SAEs arise from the *match* between the forward model prediction and sensory feedback. But he revises it to exclusively rely on forward model predictions in order to accommodate cases in which agentic experiences seem to be present despite there being no sensory feedback to match with a forward model prediction, as in the case of deafferented individuals and when one is under anaesthesia.

⁸ This reasoning seems problematic, however, since the two could be compared simply in terms of their representational content, in the way that a belief about something might be compared with a perception of that thing in terms of the respective content of each state, despite having different representational formats. I set aside this worry here, however, to focus on other aspects of Prinz's view.

particular those in which there is an error in action execution, forward model predictions will not match the sensory feedback from bodily movement. In these cases, we should expect there to be awareness of a sensory image of the intended bodily movement, and, subsequently, a perception or sensation of an erroneous bodily movement. After all, Prinz claims that the reason we are not typically aware of these sensory images is that they typically do match in content with the experience of bodily movement. One might reasonably suppose, in a similar fashion, that when they do not match, we are aware of them.

This does not seem to capture the phenomenology of action error, however. When we make such errors, we are not first presented with the conscious awareness of having successfully performed the action we are trying to perform—we are simply aware that something has gone wrong. The proponent of Prinz's view might reply that the forward model predictions are always nonconscious in these specific cases, but we would need some independent reason to think this, and to explain why they fail to become conscious in just these cases and not in others. Prinz's imagistic forward model predictions are thus not tenable candidates for SAEs.

If agentive experiences are sensory experiences, it is unlikely, given the failure of leading proposals, that they are sensory experiences in any of the familiar modalities. I move on now to consider an alternative account on which SAEs are thought to arise out of a novel, previously unrecognized sensory modality.

4.2 Against a non-reductive account of SAEs

Recently, Bayne (2011) has defended a sensory model of agentive awareness, proposing that agentive experiences exist, not in any of the standard sensory modalities, but in a novel sensory modality that is dedicated to our own agency. As he puts it:

Where in the cognitive architecture should we locate such experiences? They are not located within the central cognition, nor are they located within the systems responsible for programming and executing actions, nor are they located within the high-level reaches of any of the standard perceptual modalities. Instead, such states are the products of a dedicated perceptual system (or systems) (358).

More specifically, Bayne maintains that this sensory system is located at the site of the comparator that is also responsible for sensory attenuation—what I will call the *forward model system*. Bayne's claim is that, in cases of a match between the forward model prediction and sensory feedback, the forward model system outputs an SAE. The SAE in question is something distinct from any sensory attenuation that might subsequently follow. In this way Bayne's proposal is set apart from that of Frith and his colleagues.

There has been much debate surrounding whether the forward model system adequately accounts for all the cases in which agentive experiences putatively arise. For example, Synofzik et al. (2008) argue that it cannot adequately explain a range of cases, while Carruthers (2012) is more sympathetic, arguing instead that with suitable adjustments, it can do the explanatory work required of it. Wong (2012)

further points out that if one is too liberal with the adjustments, as may be required for sufficient explanatory power, the question arises as to whether “the term ‘comparator’ is just standing proxy for whatever mechanism is responsible for the sense of agency insofar as some feedback is involved” (50). These are important issues to settle, but I put them aside here. My focus instead will be to argue that the forward model system is not suitable for supporting SAEs at all, let alone explaining the range of cases in which they allegedly arise, and that Bayne’s proposal should therefore be rejected.

One major difficulty with the view that the forward model system is the locus of SAEs concerns the claim that the forward model system is properly construed as its own sensory modality. In order to evaluate this claim, we may appeal to characteristic features of sensory modalities, and determine whether the forward model system exhibits them. A commonly attributed hallmark of a sensory modality is that it operates by way of a sense organ over which we have some degree of control, either directly or indirectly (cf. Shoemaker 1994). When it comes to the forward model system, however, there is no candidate sense organ over which we have any degree of control. The comparison that takes place between forward model predictions and sensory feedback is not something that we can initiate, guide, or inhibit at will—either directly or indirectly. So the case for viewing the forward model system as a sensory modality seems, in this crucial respect, to break down.

Bayne (2011) anticipates this worry, and urges that the characterization of a sense organ as something over which we have some degree of control is misguided, since we do not have any control over the mechanisms responsible for proprioception, nociception, and vestibular sense either, and yet they still qualify as forms of perception. In light of this, Bayne suggests that a sense organ be construed instead as “a dedicated mechanism that takes as input raw energy of some kind and generates representations in an appropriate format, at least some of which are experiential” (370). On this characterization, the mechanisms underlying proprioception, nociception, and vestibular sense would arguably qualify as sense organs.

This does not successfully address the problem, however. First, if one accepts Bayne’s characterization of a sense organ, the forward model system still falls short. For instead of taking as input raw energy, e.g., wavelengths, as Bayne himself points out, it takes psychological representations as input, namely the forward model prediction and sensory feedback. This is rather unlike any sensory modality with which we are familiar, and stretches the notion of a sensory modality beyond recognition (see also Keeley 2002; Macpherson 2011).

Bayne (2011) contends that this difference in inputs between typical sensory modalities and the forward model system is not a problem. The forward model system, he suggests, should simply be viewed as a “non-basic” sensory modality, whereas the sensory modalities that take raw energy as input should be construed as “basic” modalities. But we have no independent reason to bifurcate sensory modalities into basic and non-basic kinds, and, furthermore, we do have reason to avoid doing so.

First, the properties of psychological representations cannot be understood in a way that is remotely analogous to the sensory properties, e.g., colors and shapes, of physical objects on which familiar sensory modalities operate. More specifically,

these latter properties stand in relations of similarity and difference to other properties in that same family. For example, red is more similar to orange than it is to blue. There are no analogous similarity and difference relations between the properties of forward model predictions.

Second, adopting Bayne's characterization of non-basic sensory modalities leads to an excessively liberal view of sensory modalities. Consider, for example, the case of arithmetic calculations. Suppose that one compares two sums with respect to which one is the greater sum, and that the output of such a comparison is a judgment to the effect that the first sum is greater than the second sum. It would seem to follow, on Bayne's characterization of sensory modalities, that this is an instance of sensory processing, and that we have a sensory modality dedicated to arithmetic. After all, the inputs are psychological representations, i.e., each calculated sum, and, if we countenance cognitive phenomenology, then they may output representations, i.e., thoughts, that are at least sometimes "experiential" in the sense of having phenomenal properties. But a view on which it follows that we have an arithmetic sense has led us far astray from any informative or familiar notion of sensory modalities.

In sum, there are significant reasons to deny that the forward model system is properly construed as a sensory modality. And if so, then there is reason to deny that it could be home to SAEs.

But perhaps one will not be moved by the above considerations. Even still, there is another angle that reveals a significant difficulty with Bayne's (2011) proposal, which is that the forward model system would be predicted to yield SAEs in cases in which it does not. Consider again the case of Mrs. GP, the individual with anarchic hand syndrome. Mrs. GP is not aware of the movements of her anarchic hand as her own actions, but they still display intact sensorimotor control. Her hand is able to locate and interact with objects in her environment without any difficulty. The movements of her arm and exercise of her muscles are initiated and guided by motor commands, and the comparison between the forward model prediction and sensory feedback results in matches for the purpose of controlling the movements of her hand. If there were mismatches yielded by this comparison, her control of these movements would be impaired at the sensorimotor level, and it is not.

If there are matches between the forward model prediction and sensory feedback, however, it is predicted, on Bayne's account, that Mrs. GP would have agentive experiences associated with the movements of her anarchic hand. Since she does not, this provides a powerful counterexample against Bayne's non-reductive version of the sensory approach.

A proponent of a sensory account of agentive experiences grounded in the comparator model may reply here that agentive experiences arise, not just out of comparisons between forward model predictions and sensory feedback, but out of comparisons between intentions and the forward model prediction as well. Indeed, this second comparison is widely posited among theorists working in this area (cf. Frith et al. 2000; Synofzik et al. 2008; Pacherie 2008). Since Mrs. GP does not have an intention guiding the actions of her anarchic hand, there would be a mismatch with respect to this second comparison, and so this would explain her lack of agentive experience.

But, even if we grant that this second comparison takes place as well, we still have no reason to view a system that takes as input intentions and forward model predictions and compares them as a sensory system. Indeed, all the reasons presented for denying that the forward model system is a sensory system will apply with equal force here. In addition, even on such a view, Mrs. GP should still have some agentic experience on the basis of the match between her forward model predictions and sensory feedback, and this does not seem to be the case. This revised version of the non-reductive account cannot therefore establish the existence of SAEs.

5 Conclusion

Agentic experiences are commonly posited in discussions of agentic awareness. But rarely is their precise nature considered. In this paper, I have argued against a particular view of such experiences, namely the view that they are sensory in form. Not only do leading arguments in favor of positing SAEs fail, but there is no sensory modality or set of sensory modalities—either familiar or novel—within which they can be plausibly located. The upshot is that, if we are to posit agentic experiences, it is more promising to look to non-sensory models to explain and describe them. And if such models are found as equally wanting as their sensory counterparts, we must altogether abandon the view that there are agentic experiences.

Acknowledgments I'm very grateful to Jacob Berger, Grace Helton, David Rosenthal, Frédérique de Vignemont, and audiences at the Graduate Center, CUNY, and NYU for helpful feedback on earlier drafts of this paper.

References

- Albers, J. (2006). *Interaction of color*. New Haven: Yale University Press.
- Armstrong, D. M. (1968). *A materialist theory of mind*. London: Routledge & Kegan Paul Ltd.
- Bayne, T. (2011). The sense of agency. In F. Macpherson (Ed.), *The senses* (pp. 355–374). Oxford: Oxford University Press.
- Bayne, T., & Pacherie, E. (2007). Narrators and comparators: The architecture of agentic self-awareness. *Synthese*, *159*, 475–491.
- Bermúdez, J. (2010). Action and awareness of agency: Comments on Christopher Frith. *Pragmatics and Cognition*, *18*(3), 584–596.
- Blakemore, S.-J., & Frith, C. (2003). Self-awareness and action. *Current Opinion in Neurobiology*, *13*, 219–224.
- Blakemore, S.-J., Frith, C. D., & Wolpert, D. M. (1999). Spatio-temporal prediction modulates the perception of self-produced stimuli. *Journal of Cognitive Neuroscience*, *11*(5), 551–559.
- Blakemore, S.-J., Goodbody, S. J., & Wolpert, D. M. (1998). Predicting the consequences of our own actions: The role of sensorimotor context estimation. *The Journal of Neuroscience*, *18*(18), 7511–7518.
- Blakemore, S.-J., Oakley, D. A., & Frith, C. D. (2003). Delusions of alien control in the normal brain. *Neuropsychologia*, *41*(8), 1058–1067.
- Blakemore, S.-J., Smith, J., Steel, R., Johnstone, C. E., & Frith, C. D. (2000a). The perception of self-produced sensory stimuli in patients with auditory hallucinations and passivity experiences: Evidence for a breakdown in self-monitoring. *Psychological Medicine*, *30*(5), 1131–1139.

- Blakemore, S.-J., Wolpert, D., & Frith, C. (2000b). Why can't you tickle yourself? *NeuroReport*, *11*(11), R11–R16.
- Block, N. (2003). Mental paint. In M. Hahn & B. Ramberg (Eds.), *Reflections and replies: Essays on the philosophy of Tyler Burge*. Cambridge, MA: MIT Press.
- Byrne, A. (2009). *Experience and Content*. *Philosophical Quarterly*, *59*, 429–451.
- Campbell, J. (1999). Schizophrenia, the space of reasons and thinking as a motor process. *The Monist*, *82*(4), 609–625.
- Carter, J. D., Bizzell, J., Kim, C., Bellion, C., Carpenter, K. L., Dichter, G., & Belger, A. (2010). Attention deficits in schizophrenia—preliminary evidence of dissociable transient and sustained deficits. *Schizophrenia Research*, *122*(1–3), 104–112.
- Carruthers, G. (2012). The case for the comparator model as an explanation of the sense of agency and its breakdowns. *Consciousness and Cognition*, *21*, 30–45.
- Chalmers, D. (2004). The representational character of experience. In B. Leiter (Ed.), *The future for philosophy* (pp. 153–181). Oxford: Oxford University Press.
- de Vignemont, F. (2011). Embodiment, ownership and disownership. *Consciousness and Cognition*, *20*, 82–93.
- de Vignemont, F., & Fournieret, P. (2004). The sense of agency: A philosophical and empirical review of the “who” system. *Consciousness and Cognition*, *13*, 1–19.
- Della Sala, S. (2005). The anarchic hand. *The Psychologist*, *18*(10), 606–609.
- Desantis, A., Weiss, C., Schutz-Bosbach, S., & Waszak, F. (2012). Believing and perceiving: Authorship belief modulates sensory attenuation. *PLoS One*, *7*(5), e37959.
- Frith, C. (1992). *The cognitive neuropsychology of schizophrenia*. Hove: Lawrence Erlbaum Associates Ltd.
- Frith, C. (2007). *Making up the mind: How the brain creates our mental world*. Oxford: Blackwell.
- Frith, C. D., Blakemore, S., & Wolpert, D. M. (2000). Explaining the symptoms of schizophrenia: Abnormalities in the awareness of action. *Brain Research Reviews*, *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.
- Gold, J. M., Fuller, R. L., Robinson, B. M., Braun, E. L., & Luck, S. J. (2007). Impaired top-down control of visual search in schizophrenia. *Schizophrenia Research*, *94*(1–3), 148–155.
- Haggard, P., Clark, S., & Kalogeras, J. (2002). Voluntary action and conscious awareness. *Nature Neuroscience*, *5*(4), 382–385.
- Harman, G. (1990). The intrinsic quality of experience. *Philosophical Perspectives*, *4*, 31–52.
- Keeley, B. L. (2002). Making sense of the senses: Individuating modalities in humans and other animals. *Journal of Philosophy*, *99*(1), 5–28.
- Lindner, A., Thier, P., Kircher, T. T., Haarmeier, T., & Leube, D. T. (2005). Disorders of agency in schizophrenia correlate with an inability to compensate for the sensory consequences of actions. *Current Biology*, *15*(12), 1119–1124.
- MacPherson, F. (2011). Individuating the senses. In F. Macpherson (Ed.), *The senses* (pp. 3–43). Oxford: Oxford University Press.
- MacPherson, F. (2012). Cognitive penetration of colour experience: Rethinking the issue in light of an indirect mechanism. *Philosophy and Phenomenological Research*, *84*(1), 24–62.
- 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.
- Mellors, C. S. (1970). First-rank symptoms of schizophrenia. *British Journal of Psychiatry*, *117*, 15–23.
- Metzinger, T. (2006). Conscious volition and mental representation: Toward a more fine-grained analysis. In N. Sebanz & W. Prinz (Eds.), *Disorders of volition* (pp. 19–48). Cambridge, MA: Bradford Books.
- Mylopoulos, M. I. (2011). Why reject a sensory imagery theory of control consciousness? *Topics in Cognitive Science*, *3*(2), 268–272.
- Mylopoulos, M. (2012). Evaluating the case for the low-level approach to agentive awareness. *Philosophical Topics*, *40*(2), 103–127.
- Pacherie, E. (2008). The phenomenology of action: A conceptual framework. *Cognition*, *107*, 179–217.
- Peacocke, C. (2003). Awareness, ownership, and knowledge. In J. Roessler & N. Eilan (Eds.), *Agency and self-awareness: Issues in philosophy and psychology* (pp. 94–110). Oxford: Oxford University Press.

- Penfield, W. (1975). *The mystery of the mind: A critical study of consciousness and the human brain*. Princeton, NJ: Princeton University Press.
- 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.
- Prinz, J. J. (2012). *The conscious brain: How attention engenders experience*. New York, NY: Oxford University Press.
- Rosenthal, D. M. (2004). Varieties of higher-order theory. In R. J. Gennaro (Ed.), *Higher-order theories of consciousness: An anthology* (pp. 17–44). Philadelphia, PA: John Benjamins B. V.
- Rosenthal, D. M. (2005). *Consciousness and mind*. New York: Oxford University Press.
- Searle, J. R. (1983). *Intentionality: An essay in the philosophy of mind*. Cambridge, MA: Cambridge University Press.
- Shergill, S. S., Samson, G., Bays, P. M., Frith, C. D., & Wolpert, D. M. (2005). Evidence for sensory prediction deficits in schizophrenia. *American Journal of Psychiatry*, 162(12), 2384–2386.
- Shoemaker, S. (1994). Self-knowledge and “inner sense”: Lecture I: The object perception model. *Philosophy and Phenomenological Research*, 54(2), 249–269.
- 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., & Newen, A. (2008). Beyond the comparator model: A multifactorial two-step account of agency. *Consciousness and Cognition*, 17(1), 219–239.
- Tye, M. (2000). *Consciousness, color, and content*. Cambridge, MA: MIT Press.
- Tye, M. (2003). *Consciousness and persons: Unity and identity*. Cambridge, MA: MIT Press.
- Wegner, D. (2002). *The illusion of conscious will*. Cambridge, MA: Bradford Books.
- Weiller, C., Juptner, M., Fellows, S., Rijntjes, M., Leonhardt, G., Kiebel, S., et al. (1996). Brain representation of active and passive movements. *Neuroimage*, 4(2), 105–110.
- Wolpert, D. M. (1997). Computational approaches to motor control. *Trends in Cognitive Sciences*, 1(6), 209–216.
- Wong, H. Y. (2012). A measure of my agency? *Consciousness and Cognition*, 21(1), 48–51.