

The Experience of Agency

Myrto Mylopoulos and Joshua Shepherd

In this chapter we reflect on questions about the nature and sources of agentic phenomenology – that is, the set of those experience-types associated with exercises of agency, and paradigmatically with intentional actions. Our discussion begins with pioneering work in psychology and neuroscience that dates to the early 80s (section 1). As we will see, much of the current work on agentic phenomenology in both psychology and philosophy draws motivation from this work, and the questions it raises.¹ After discussing empirical work relevant to agentic phenomenology, we turn to consideration of its nature. We cover questions about the scope of agentic phenomenology, about its relationship to other types of experiences (section 2.1), about the best way to characterize aspects of agentic phenomenology, and about the function of various types of agentic experience (section 2.2).

1 The science of agentic phenomenology

1.1 Consciously deciding

In a now-classic study, Libet, et al. (1983) asked six participants each to perform a “quick, abrupt flexion of the fingers/and or wrist” (625) with their right hand, at a time of their own choosing. They were instructed not to decide in advance on a time at which to perform the act, but rather “to let the urge to act appear on its own at any time without any preplanning or concentration on when to act” (625). Participants were seated facing a specialized clock, around which a dot would revolve every 2.56 seconds. They were instructed to report the time at which they first became aware of deciding or having the “urge” to move based on the position of the dot on the clock (the “W judgment”). They were also asked to report the time at which they were aware that they had “actually moved” (627) (the “M judgment”). During the task, they were hooked up to an electroencephalogram (EEG) that measured their brain activity, as well as an electromyogram (EMG) that measured their muscle activity.

The most striking result that Libet and colleagues uncovered was this.² A slow build-up of neural activity preceding the onset of spontaneous movements (known as the ‘Readiness Potential or ‘RP’), and which has been widely associated with the neural signature of action initiation (though see Schurger, Sitt, & Dehane 2012; Schurger, Mylopoulos, and Rosenthal

¹ Certainly one also finds motivation stemming from work on agentic phenomenology in the history of psychology and philosophy. Regarding psychology, see de Biran (1812) and James (1880). Regarding philosophy, see, e.g., Brentano (1874), Sartre (1969), Merleau-Ponty (1996), and Ricoeur (1966). For an interesting recent discussion of Hume’s phenomenology of agency, see Wood (2014).

² Another notable result from the Libet study is this. Participants reported awareness of acting on average 86 ms prior to the onset of bodily movement as measured by EMG. This might be viewed as *prima facie* support for accounts that take agentic phenomenology to arise prior to any sensory feedback from bodily movement, thus providing motivation for cognitive accounts of agentic phenomenology based on intentions, tryings, or doxastic states. Others are skeptical that the results can be interpreted in this way, arguing that the subjective reports provided by participants are not conclusive (e.g., Carruthers 2012).

2016; and Mele 2010 for competing interpretations), precedes the time at which participants report being aware of an urge or decision to act by an average of approximately 350 ms. (The reported time of the urge or decision itself precedes the onset of the action by about 200 ms.) As Libet and colleagues put it, “the brain evidently ‘decides’ to initiate or, at the least, prepare to initiate the act at a time before there is any reportable subjective awareness that such a decision has taken place” (Libet et al. 1983: 640). This result, and especially the interpretation that Libet and colleagues offer, has garnered an astounding amount of scholarly and popular attention. This is largely due to assumed implications regarding free will. But it is important to note that these implications depend on assumptions regarding the phenomenology of deciding. For the natural thought is that if ‘the brain decides’ such that the phenomenology of deciding is illusory, the person must not be in control, and must not therefore be free.

Libet’s work, coupled with the many studies that followed from Libet’s original one, has prompted a surge of interest in the phenomenology of deciding (see, e.g., Carruthers 2007, Mele 2009, Shepherd 2013). While some have broadly taken Libet’s side, arguing that the phenomenology of deciding is in some sense illusory (e.g., Carruthers 2007), most have argued that Libet’s results at best demonstrate the existence of early bias in decision-making processes (e.g., Mele 2009). How best to understand the processes underlying decision-making, and the relationship of the phenomenology of deciding to those processes, remains undetermined (though see Bear and Bloom 2016 for an interesting recent study).

1.2 The feeling of doing

A second relevant research program was led primarily by the psychologist Daniel Wegner. Wegner focused on the psychological mechanisms giving rise to what he called the ‘feeling of doing’. Wegner (2002) argues on the basis of folk theoretical observations, as well as work demonstrating that priming for certain action outcomes influences the degree to which agents report feeling in control of their behavior (Wegner and Wheatley 1999), that our experience of consciously willing our actions is illusory. For Wegner, the underlying causes of our behaviour are unconscious processes in the brain, and not our conscious intentions themselves. He develops a view on which we infer that our intentions cause our actions when we are aware of them occurring prior to the action, they are consistent with the action, and there are no salient alternative causes. In this way, “The experience of will may be a result of the same mental processes that people use in the perception of causality more generally” (Wegner 2004: 654).

Regarding the influence of Wegner’s work, three similarities with Libet deserve mention. First, as with Libet, much of the scholarly and popular attention devoted to Wegner’s results has centered on assumed implications for free will. Second, these implications depend on assumptions regarding agentive phenomenology – in this case, the feeling of doing. As with Libet, the evident importance of agentive phenomenology to questions of broad societal interest motivated renewed interest in agentive phenomenology. Wegner himself characterized the ‘feeling of doing’ in multiple ways (see Nahmias 2005 for discussion). In response to Wegner’s work, thinkers interested in agentive phenomenology have thus been forced to consider just what we might mean by speaking of a feeling of doing, and how best to characterize the kinds of experiences that constitute it. Third, as with Libet, the majority of philosophers have argued that these results, while interesting, pose little threat to free will (see especially Mele 2009, Bayne 2011c, Shepherd 2015a).

1.3 The comparator model and agentive phenomenology

A third line of work attempts to understand agentic phenomenology by anchoring it to the states and processes underlying sensorimotor control. This approach leans heavily on a popular neurocomputational model of how such control functions known as the Comparator Model (CM). According to this model, several computational mechanisms work together to anticipate upcoming segments of action execution, to check for possible mistakes, and to correct for mistakes in real time. One mechanism, called the inverse model, computes motor commands – representations that specify the fine-grained aspects of bodily movement, such as grip aperture and grip force – on the basis of the current state of the body in order to satisfy a given goal state. Another mechanism, called the forward model, is thought to take as input a copy of the motor command, known as the efference copy, as well as a representation of the current state of the body, and output a prediction of the sensory consequences that executing that command would entail. In addition, a series of comparisons involving the goal state, the forward model prediction, and the sensory feedback from the bodily movement are posited in order to explain various control operations.

First, there is a comparison between the forward model prediction and the sensory feedback. Second, just prior to movement, there is a comparison between the forward model prediction and the goal state. If a mismatch results, then an error signal is sent to the inverse model, and a new motor command is computed. Lastly, a third comparison between the goal state, i.e., a proximal intention, and sensory feedback is posited. This comparison plays a role in motor learning, since it allows the training up of the inverse model, but also ongoing correction and adjustment of the action as it unfolds, if needed (Blakemore and Frith 2003; Blakemore, Frith, and Wolpert 1999; Blakemore, Oakley, and Frith 2003).

Agentic phenomenology is typically thought to arise as a result of a match between the forward model prediction and sensory feedback from bodily movement (Bayne and Pacherie 2007; Pacherie 2008; Bayne 2011b).³ As Bayne (2011b) describes it, “when there is a match between predicted and actual state, the comparator sends a signal to the effect that the sensory changes are self-generated; when there is no match (or an insufficiently robust match), sensory changes are coded as externally caused” (495).

Motivation for this model arises in part from the observation that disruptions in agentic phenomenology seem to correlate with deficits in generating accurate forward model predictions. Perhaps one of the most celebrated findings in support of this view is that, while healthy individuals experience sensory attenuation in relation to a self-generated tactile stimulus, but not an externally generated tactile stimulus (Blakemore et al. 2003), schizophrenic individuals with delusions of control do not experience such attenuation to the same degree (Blakemore, Oakley and Frith 2003). The thought is that the mechanism behind sensory attenuation is forward model prediction—when a sensory consequence is predicted by the forward model it is “cancelled out”—and schizophrenic individuals are impaired in their ability to form just such states, leading to a failure of attenuation, and thereby a lack of agentic phenomenology.

³ Some theorists propose a division of phenomenal labour, according to which the comparison between the goal state and the forward model prediction underlies one’s sense of *initiating* an action and the comparison between the forward model prediction and sensory feedback underlies one’s sense of *control* over an action (e.g., Synofzik, Vosgerau, & Newen 2008). But the main evidence appealed to in support of positing the former comparison is not compelling. It is that a number of studies have found that people tend to report anticipatory awareness of acting—that is, they report having started acting before any sensory feedback from the action is available (Lau, Rogers, Haggard, & Passingham 2004; Libet, Gleason, Wright, & Pearl 1983). Such reports, however, could just as well be based on intentions or decisions to start acting, rather than the output of any comparison involving a forward model prediction.

Another set of widely-cited studies in support of the CM account of agentic phenomenology uses action-monitoring paradigms. In these paradigms, participants are tasked with performing a simple act (e.g., moving a joystick, reaching towards a target), and presented with visual feedback that is distorted to varying degrees relative to their movement (e.g., Daprati et al. 1997; Farrer et al. 2003; Fournier and Jeannerod 1998). Participants are then asked whether it is their movement that they see, or one caused by an external source. The thought here is that participants base their judgements on the comparison between the forward model prediction and the visual feedback they receive. Again, schizophrenic individuals with delusions of control tend to perform poorly on this task compared to healthy controls and schizophrenic individuals without such delusions (Daprati et al. 1997). Moreover, the accuracy of their performance even correlates with the severity of their symptoms (Synofzik et al. 2010).

Here we suggest, in agreement with others (e.g., Synofzik et al. 2008), that though the CM may be a suitable model for explaining sensorimotor control, it faces some difficulties when it comes to explaining agentic phenomenology. Concerning the sensory attenuation results, it has been found that sensory attenuation of an auditory tone following a voluntary action is modulated by the presence or absence of a belief that one's action is the cause of the tone. But if so, then it would seem that forward model prediction, which is present across conditions, cannot be the source of the attenuation. So insofar as sensory attenuation is linked to agentic phenomenology (more on this below), it does not seem that the CM can explain it. Other results that are linked to both forward model predictions and agentic phenomenology, such as the "intentional binding effect" (Haggard, Clark, and Kalogeras 2002), i.e., the finding that the effects of one's actions are subjectively perceived as occurring earlier in time than the effects of one's involuntary movements, can also arguably be divorced from the CM in this way (see Mylopoulos 2012).

Second, it has been pointed out that without case-by-case adjustments to how different states (e.g., visual vs. proprioceptive feedback) are weighed in the comparison, the CM has difficulty consistently accounting for a range of experimental results (Carruthers 2012). For example, deafferented individuals who lack proprioceptive feedback arguably experience agentic phenomenology in virtue of a match between visual feedback from their bodily movements and forward model predictions. But in other studies, proprioception seems to play a key role in agentic experience (Balslev, Cole, and Miall 2007).

Such difficulties have led some theorists to assert that "there are many cases where the comparator output is neither a sufficient nor a necessary condition for the [feeling of agency]" (Synofzik et al. 2008: 226). An alternative model that has been offered is the so-called multifactorial weighting model, on which a range of action-related cues such as efference copy, forward model prediction, and multimodal sensory feedback, are used as inputs in a multifactorial weighting process that determines agency according to a Bayesian learning framework (Synofzik et al. 2008)

While this model may hold promise, at this point it is not yet known how to determine in a principled way what the appropriate weights are for given cases. As such, this model does not currently enable predictions for or explanations of when or why agentic phenomenology arises, and may even be unfalsifiable (see Carruthers 2012 for discussion on this point).

Third, it is remarkable that the CM approach is such that the subpersonal states and processes that purportedly give rise to agentic phenomenology operate to some degree independently of the agent's intentions, as well as action-related beliefs and perceptions. This relegates such states to the backseat in driving agentic phenomenology. But it seems

unlikely that, central as they are to intentional control, these states would not make a more direct contribution to our experiences of acting. Accounts of agentive experience pitched in terms of these personal-level states (e.g., Wegner 2002) or a mix of personal-level states and the sub-personal states of the CM (e.g., Pacherie 2008) may therefore be preferable to those that focus exclusively on the CM.

Finally, we emphasize that it is not at all clear *which* aspects of agentive phenomenology the studies and paradigms that are frequently appealed to in support of the CM are actually measuring. For instance, in the case of the intentional binding effect, only the participants' subjective timing judgements relating to the consequences of their bodily movements are explicitly probed, leaving it open as to whether the effect also indicates a sense of mineness, a sense of control, or some other aspect of agentive phenomenology. Similarly, when it comes to the action-monitoring paradigm, though it is often claimed to examine the 'sense of agency' (more on various usages of this term later in the discussion) it seems primarily to measure action perception, that is, whether one's proprioceptive feedback matches with the visual feedback one receives. Insofar as 'the sense of agency' is thought to refer to a sense of executing one's intention, or a sense of purposiveness, it is not being directly measured by this study. (Consider that the task could be performed with entirely passive movements as well.) We attempt to clear up some of these issues in section 2.2.

1.4 Metacognition and agentive phenomenology

We turn now to consider a fourth line of work gaining traction lately, suggesting that aspects of agentive phenomenology can be explained by appeal to metacognitive states and processes (e.g., Wenke, Fleming, Haggard 2010; Metcalfe et al. 2013; Chambon et al. 2014; Carruthers forthcoming). In particular, some have recently suggested that agentive phenomenology reflects the metacognitive monitoring of aspects of our action-producing mechanisms as they are being engaged (see also section 2.2.5 for discussion of a metacognitive account of mental agency).

One attractive feature of this view is that it can plausibly account for the subjective character of agentive phenomenology. After all, many metacognitive states are construed as subjective feelings, such as the "feeling of knowing" (Koriat 2000). Moreover, these are hardly robust experiences like the feeling of pain or a typical visual experience. It is not unreasonable, then, that agentive phenomenology, which is often itself characterized as "thin" – perhaps an indication that its subjective character is difficult to capture in virtue of being typically located at the margins rather than the center of one's total experience at a time – might be assimilated to this class of feelings as well.

Though this view holds considerable appeal, a central question concerns the aspects of the action-producing mechanisms to which agentive phenomenology is sensitive. Here we examine the recent proposal that agentive phenomenology is the result of the metacognitive monitoring of "fluency" signals generated during the process of action selection (Wenke, Fleming, and Haggard 2010; Chambon et al. 2014).

Wenke et al. explore this possibility with an elegant study asking participants to perform a left or right keypress in response to a left- or right-facing target arrow. Prior to the target arrows, participants are presented with a masked prime arrow that is either compatible or incompatible with the target. Once they press the key, a colour patch appears after a certain delay interval. The colour of the patch corresponds to whether the target arrow is compatible or incompatible with the masked prime. At the end of each block, participants are asked to rate how much control they felt they had over the specific colours that appeared. Participants reported feeling more control over prime-compatible colours, i.e.,

those occurring after actions that were primed and therefore more easily selected, versus prime-incompatible colours, despite the predictability of the colours being the same—namely zero.⁴

Some argue on the basis of such results that (i) agentic phenomenology relies on the evaluation and monitoring of action selection processes prior to action execution, and (ii) agentic phenomenology should thus be construed as a form of metacognition.

In particular, it has been suggested (Chambon et al. 2014) that fluent action selection is a reliable predictor of control over action *effects*, and in this way the monitoring of fluency signals can be construed as a metacognitive process insofar as it serves as a capacity “through which an operating cognitive subsystem [in this case the action-producing system] is evaluated or represented in a context-sensitive way” (Proust 2013: 13). Here we’d have an analogue of the “feeling of knowing”, which evaluates the potential success of mnemonic retrieval operations, in the “feeling of doing”, which evaluates the potential success of action-production processes.

While a proposal along these lines may ultimately prove fruitful, important questions remain to be settled. For one, though the most widely-discussed empirical studies in this area only explicitly ask about participants’ feelings of control over the effects of their actions, it is reasonable to suppose that the monitoring of fluency signals contributes to a sense of control over the executed actions themselves, and not just the effects that these paradigms happen to focus upon. But to our knowledge this possibility has yet to be directly examined. Second, fluency is not the only candidate for what is being monitored here. Another option is that it is the extent to which the executed action—again, rather than its effect—is predicted by one’s action producing mechanism, with the subliminal prime serving as a a predictive representation. Third, there is a question regarding whether the metacognitive monitoring being posited involves an explicit self-attribution of agency via a metarepresentation or, rather, is simply sensitive to fluency cues in some way that falls short of explicit self-attribution.

On this last point, indeed, it may be that these different forms of metacognition that are potentially at play here would respectively track different aspects of agentic phenomenology, some requiring explicit self-attribution and some not. We further stress that the terms ‘sense of agency’ and ‘sense of control’ are often used interchangeably in describing what is being examined in this work, and it’s sometimes unclear what exactly theorists have in mind when they use these terms and whether or not they are being consistent in their usage. Without more precision here, it is difficult to evaluate what specific contributions are being made to the enterprise of understanding agentic phenomenology. In section 2, we reflect on ways that philosophy has attempted to contribute to such precision, and suggest some further refinements of our own.

1.5 Summary

Work in psychology and neuroscience has been responsible for increased interest in the sources of agentic phenomenology. But as we have seen, it is debatable just what aspects of agentic phenomenology this work seeks to, and in fact does, elucidate. We have noted

⁴ In a subsequent study, Chambon et al. (2012) found that participants reported feeling more in control of prime-compatible colours than prime-incompatible colours even when reaction times were slower for the former, thus dissociating the monitoring of fluency signals from performance monitoring that uses reaction times as a cue to performance.

issues surrounding Libet's characterization of the phenomenology of deciding, Wegner's characterization of the feeling of doing, and attempts to explain the agentic phenomenology via the comparator model and metacognitive processes. These issues are not insurmountable, but they do suggest that greater clarity is needed regarding the explananda. Just what is the structure of this thing, agentic phenomenology, the sources of which science is attempting to reveal?

2 The philosophy of agentic phenomenology

Two broad issues have occupied philosophical attention to agentic phenomenology. The first issue has to do with the structure of consciousness in general, and in particular with the relationship between agentic phenomenology and the rest of conscious mental life. Is agentic phenomenology proprietarily agentic, i.e., not reducible to perceptual, emotional, or cognitive phenomenology? The second issue has to do with the structure of agentic phenomenology more specifically. What kinds of experience-types qualify as agentic, how ought we to taxonomize them, and how do they relate to each other? We discuss both issues in turn.

2.1 *Is any phenomenology proprietarily agentic?*

Consider this sequence of activity, as described by an amateur skier:

To my alarm, I recognized that the control that I had over the direction of my skis, never anything near total, was now nothing more than feeble. Each turn required immense concentration and exhausting effort . . . I became aware that I was shaking not simply through physical exhaustion but also because I was beginning to get quite frightened . . . In the half-light, the surface of the mountain appeared to be tablecloth smooth but I knew this to be nothing more than a mirage. In the daylight, every contour would be clearly visible, but now I had to try and relax enough to ride any invisible bump, hillock or mogul. Leaning forward against all my natural instincts, I bent my knees a little further trying to absorb any nasty shocks with my aching legs. Momentarily, I thought that I had skied off the edge of a cliff as my internal organs, operating on some different system of gravity to the rest of my leaden body, rose in my stomach. Several minutes later, or so it seemed, my skies came back to earth . . . Feeling delighted that I was still upright, I thumped my ski sticks back into the freezing surface. As I did so, I suddenly sensed, to my alarm, that I was sliding rapidly backwards. (Randall 2009: 5-6)

Is this an episode of *agentic* phenomenology? That depends on what it means to call a bit of phenomenology agentic. We might gloss agentic phenomenology as the conscious experiences surrounding (both bodily and mental) action. But this is a fairly weak gloss. Perceptual, emotional, and cognitive experiences (among others) surround action, and figure in the above description. Is there anything special or unique in virtue of which a bit of phenomenology might be considered agentic?

Drawing on Uriah Kriegel's recent discussion (Kriegel 2015: 73), here is one way to approach this question. Let us say that a bit of phenomenology can be considered obviously agentic if it instantiates a phenomenal property that is proprietarily agentic. And let us say that a phenomenal property is proprietarily agentic if it is

instantiated by a state or process central to some actional episodes, and not instantiated by states and processes that occur outside of actional episodes.

Why believe that any phenomenal property is proprietarily agentive in this sense? One reason would have to be based on introspecting one's conscious experiences. Accordingly, philosophers who endorse proprietarily agentive phenomenology advance largely phenomenological arguments. Horgan, Tienson and Graham (2003) and Horgan (2007, 2011) develop contrast cases based on observations about what the phenomenology of agency is not like.

[I]t is certainly not like this: first experiencing an occurrent wish for your right hand to rise and your fingers to move into clenched position, and then passively experiencing your hand and fingers moving in just that way. Such phenomenal character might be called the phenomenology of fortuitously appropriate bodily motion. It would be very strange indeed, and very alien. Nor is the actional phenomenological character of the experience like this: first experiencing an occurrent wish for your right hand to rise and your fingers to move into clenched position, and then passively experiencing a causal process consisting of this wish's causing your hand to rise and your fingers to move into clenched position. Such phenomenal character might be called the passive phenomenology of psychological state-causation of bodily motion. (Horgan 2007: 8)

Horgan and colleagues take it, then, that what the phenomenology of agency is like cannot be characterized without reference to a proprietarily agentive property. They call this property self-as-source – 'the what-it's like of self as source of the motion' (2007: 8).

Uriah Kriegel (2015) offers a different characterization of proprietarily agentive phenomenal properties. Rather than self-as-source, Kriegel emphasizes the phenomenology of deciding, which involves a 'felt pull' to action, and of trying, which satisfies the tension inherent in decision's felt pull and involves the experience of 'mobilizing force in the face of resistance' (90). Kriegel considers a number of proposals that would reduce such phenomenal properties to perceptual or emotional properties, and argues they are all inadequate. Perhaps the most plausible such proposal is due to William James, and emphasizes anticipatory imagery (for a similar proposal, cf. Prinz 2007). As Kriegel notes, 'on this view, the key element for capturing the conative dimension of the experience of clenching one's fist is the feel of imaginatively anticipating one's fist muscles contracting' (80). Against this proposal, Kriegel argues, first, that it gets the timing of the phenomenology wrong, placing it before the doing takes place. Second, it 'seems false to our experience.' Kriegel elaborates:

We experience a representation of the act to follow, but also of the act following, and following because we *make* it follow. That is, we experience not only an anticipation of the act, but also the causing of the act in real time. (80)

Although introspection has to play some role in an argument for proprietarily agentic phenomenology, it is important to note that empirical considerations can be brought to bear as well. In a recent paper Shepherd (2016) offers an empirically based argument against reductive proposals like James's. Shepherd observes that the best empirical reason to adopt a proposal like James's requires a close connection between phenomenology and the kinds of computational models posited to explain action control. Our best such models give an important functional role to states of anticipatory imagery. As we discussed earlier in the context of forward model predictions, such states are used to compare anticipated with actual perceptual feedback, in order to identify and correct errors in real time. But is there good reason to tie phenomenology to such states (and not others)?

Against such a move, Shepherd (2016) appeals to studies involving temporary paralysis. In these studies, participants try to perform various actions with various body parts. They report strong experiences of trying to move, but very little in the way of anticipatory imagery. Shepherd comments:

The experiences of trying reported in these [experiments] cannot be explained by matches between predicted and actual feedback, since there was no actual feedback. Nor can they be explained by internally simulated feedback, or anticipations of feedback. For such feedback, if centrally involved in the experiences of trying, should produce experiences of the feedback simulated or anticipated. That is, such experiences should be (presumably, sensory) experiences of things happening at the relevant bodily sites. But . . . very little sensory experience was reported. The experiences reported were largely directive in nature, concerning the direction of effort to the body parts. And . . . experiences of trying were clearly distinct from (though seemingly causally linked with) the sensory experience of body parts moving. (425-26)

Mylopoulos (2015) advances a different set of empirical considerations against views that would attempt to explain agentic phenomenology in terms of purely sensory phenomenology. Mylopoulos notes that such views come in reductive and non-reductive varieties. Reductive views attempt to identify agentic phenomenology with the phenomenology of familiar sensory modalities – most plausibly, proprioception, vision, or some combination. But such proposals have problems accommodating the full range of empirical data on action and experience. For example, stimulation of the motor cortex produces movement (as well as proprioceptive and visual experience of the body movement) but not – according to those whose motor cortex has been stimulated – any agentic phenomenology. Indeed, as Mylopoulos notes, it is difficult to see how a reductive proposal could account for the distinction we readily make between action and passive movement, for there seems to be no set of sensory qualities that mark the difference between 'actions as against non-actions' (2015: 774; see also Bermúdez 2010 for further discussion of this argument).

Non-reductive attempts at redescription must identify agentic phenomenology with a novel sensory modality. Bayne (2011b) makes a case for such a modality, connecting it to the Comparator Model of action control discussed above. In response, Mylopoulos (2015) observes that it is not clear why we should consider phenomenology attached to this model *sensory*. Further, Bayne's proposal makes failed predictions. Patients with anarchic hand syndrome display sophisticated

sensorimotor control while denying that the movements of their hands are their own actions. But if agentive phenomenology was closely tied to the computational mechanisms responsible for sensorimotor control, anarchic hand patients should have agentive phenomenology. Mylopoulos notes that those sympathetic to Bayne's proposal might wish to invoke some additional element at some point – perhaps the presence of an intention in addition to low-level sensorimotor control. Doing so, however, further undermines the thought that the relevant phenomenology is sensory.

At the very least, then, it appears that the skeptic about proprietary agentive phenomenology faces serious problems. The skeptical view appears to lack empirical motivation. Further, the skeptical view is phenomenologically inadequate. The case for proprietarily agentive phenomenology appears strong.

2.2 Describing agentive phenomenology

Agentive phenomenology is complex. In addition to involving proprietarily agentive aspects, agentive phenomenology often involves a wide range of perceptual, emotional, and cognitive (as well as, arguably, metacognitive: see section 1.4) experiences. Further, agentive phenomenology involves these experience-types in dynamic relationships that differ depending upon the agent's action-type, skill-level, and circumstances. Taxonomizing the relevant experience-types, and explaining their dynamic relationships with one another, is thus a difficult task. Unsurprisingly, those writers who have attempted the task have often used idiosyncratic terminology, and have often placed emphasis in different places.

Agentive phenomenology is often referred to simply as 'the sense of agency'. A glance at the literature reveals the many different aspects of agentive phenomenology that this term has been used to capture. Some authors use it to emphasize the subjective feeling of control, as when Haggard & Chambon (2012) describe it as "the experience of controlling one's own actions, and, through them, events in the outside world" (R390). Other times, the focus is on self-attributing an action, as when Jeannerod (2007) writes that it is "the ability to identify oneself as the agent of a behavior or a thought" (64) and Gallagher (2000) glosses it as "[t]he sense that I am the one who is causing or generating an action" (15). Some prefer to focus on authorship or ownership, as when Marcel (2003) depicts it as "[a] sense of oneself as an actor or a sense that actions are one's own" (54), and Pacherie (2007) defines it as "the sense the agent has that he or she is the author of that action" (2). Finally one also finds the term being used to single out the sense of initiating or causing an action, as when Synofzik, Vosgerau, & Newen (2008) state that it is "the registration that we are the initiators of our own actions" (219) and Sato (2009) characterizes it as the "sense that one is causing an action" (74). At times there is also some focus not just on causing bodily actions, but the consequences of such actions, as when Blakemore and Frith (2003) portray it as "the feeling that we cause movements and their consequences" (220). In light of all the various ways that 'the sense of agency' has been used, we suggest retiring it in favour of more precise terminology.

To that end, in this section we propose an opinionated taxonomy that does as much justice as possible to what others have said. Our aim is to identify areas of implicit

agreement and disagreement, and to make progress towards a more refined understanding of the descriptive shape of agentive phenomenology.⁵

In our view, reflection on agentive phenomenology – and reflection on what others have said about it – suggests six somewhat dissociable clusters. We will call these clusters the phenomenology of purposiveness (2.2.1), the phenomenology of mineness (2.2.2), the phenomenology of execution (2.2.3), the phenomenology of action perception (2.2.4), the phenomenology of action assessment (2.2.5), and the phenomenology of free will (2.2.6). Understanding the nature of these clusters, and how they conjoin and overlap in full-blown agentive phenomenology, is the central descriptive task facing the phenomenologist.

2.2.1 Purposiveness

A number of authors emphasize the purposiveness of agentive phenomenology. For example, David Hume (2000) speaks of “the internal impression we feel and are conscious of, when we knowingly give rise to any new motion of our body or new perception of our mind” (257). His language suggests that he regards the ‘giving rise to’ as fundamentally purposive.

Paul Ricoeur seems to echo the same idea. Contrasting the experience of willing with the experience of the body in action, Ricoeur writes “the personal body presents itself as body-moved-by-a-willing, that is, as the terminus of a movement which *comes down* from the ‘I’ to its mass” (1966: 220).

Both Carl Ginet and Brian O’Shaughnessy invoke cases without perceptual feedback to assert a similar view.

It could seem to me that I voluntarily exert a force forward with my arm without at the same time its seeming to me that I feel the exertion happening: the arm feels kinesthetically anesthetized. (Ginet 1990: 28)

Consider now the event of trying to move an arm whose psychologicality was revealed nakedly to view because of the complete absence of kinaesthetic experience . . . where was that psychological event, of which one was immediately aware, experienced as taking place? One is perhaps inclined to say one experienced it as being at the felt location of the arm. However, it seems to me that to do so would be to confuse the site of the target-object for the will, with the event of willing that is directed to that place. (O’Shaughnessy 2003: 351)

More recently, Uriah Kriegel characterizes the experience of trying – which as mentioned is a part of what Kriegel regards as the core of agentive phenomenology – in purposive terms, as ‘a nonsensory analog of innervation (a feeling of a kind of

⁵ Further questions include the ways these aspects of agentive phenomenology appear in various domains of action. For example, what is the best characterization of the phenomenology of expert (or skilled) action (see Christensen et al. 2016, Dow forthcoming, Shepherd 2015b)? What is the phenomenology of joint action like, and how does it differ from that of individual action (Pacherie 2012)? These questions deserve separate treatment, and as such will not occupy our attention here.

nonsensical current traveling from will to muscle)' (2015: 95). And Shepherd has characterized the experience of trying as follows:

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. It is this fundamentally directive character that marks the experience out as an experience of trying. (2016: 421)

Regarding purposive phenomenology, a number of questions remain open. It is natural to consider purposiveness as that aspect of agentive phenomenology that is most clearly proprietary in the sense describe above. But is this right? What is the relationship between purposiveness and intention? What is the relationship between purposiveness and the mechanisms that enable action control? We cannot discuss these questions here – we mention them as bookmarks for future research.

A further interesting question concerns the relationship between purposiveness and other aspects of agentive phenomenology – in particular, what we describe below as mineness phenomenology. Many of the passages above invoke the 'I' or the agent or some central place from which the experience emanates in describing the structure of purposive phenomenology. How are we to understand such language?

2.2.2 *Mineness*

It is often asserted that in consciously acting, agents experience their activity as in some sense their own. Horgan, Tienson, and Graham (2003) call this feature self-as-source phenomenology. According to them, '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 else as being transeuntly caused by your own mental states' (329). Similarly, Bayne and Levy (2006) emphasize what they call the experience of authorship, maintaining that 'it is not unlikely that the experience of authorship is essential to the experience of agency – that to experience a movement as one's own action necessarily involves an experience of oneself as the author of the movement' (56).

Beyond such initial assertions, how ought we to understand the phenomenology of mineness? One attractive option is to understand the phenomenology of mineness as implicit in the phenomenology of purposiveness. Consider the following characterization of the experience of trying.

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. It is this fundamentally directive character that marks the experience out as an experience of trying. (Shepherd 2016: 421)

On the option under consideration, the experience as of mandating or directing bodily activity is at once purposive and implicitly agent-involving. That is to say, there need be no experience of a substantive self involved – the directive aspect of the phenomenology suffices to mark out the experience as both purposive, and as the agent’s own.

A second option maintains that mineness and purposiveness are distinct, and that mineness is a primitive feature of agentive phenomenology. The italics in the following passage suggest that this is how Terry Horgan views mineness phenomenology.

How . . . should one characterize the actional phenomenal dimension of the act of raising one’s hand and clenching one’s fingers, given that it is not the phenomenology of fortuitously appropriate bodily motion and it also is not the passive phenomenology of psychological event-causation of bodily motion? Well, it is the what-it’s-like of *self as source* of the motion. 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*. (2007: 8)

A third option is to understand mineness phenomenology as a kind of cognitive or thought-like phenomenology, attached to a belief, knowledge-state, or thought to the effect that one is producing the action in question. In a recent paper Mylopoulos defends this option at length, arguing that agentive thoughts that one is A-ing can be viewed as anchoring mineness phenomenology in virtue of deploying the essential indexical.

It is worth noting that in addition to the intrinsic interest one might take in describing mineness phenomenology, one’s descriptive account might have philosophical ramifications elsewhere. For example, agent causalists – those who hold that actions are caused by agents, as opposed to events or facts – might appeal to the phenomenology of mineness as support for their view. Whether such a move is sound will depend in part upon the nature of the phenomenology (see Bayne 2008 for discussion).

2.2.3 Execution

A third aspect of agentive phenomenology, closely related to purposiveness, is that of *execution*. Purposiveness involves aiming or directing one’s efforts towards a certain outcome. But the phenomenology of agency often involves the sense that one is *doing* what one intends to do—not merely the sense of striving towards some goal, but of successfully achieving it.

Empirical work, too, suggests a dissociation between the subjective awareness of purposiveness and of action execution. In Libet’s classic subjective timing studies (1983, 1985), participants reported a feeling of “deciding” or having an “urge” to act that clearly preceded any awareness of performing the action in question (see footnote 2).

Having the sense that one is executing one’s intended action is often intimately tied to an awareness of what one’s body is doing and an evaluation of how *well* one is

performing the action in question. We'll discuss each of these further facets of agentic phenomenology in what follows, but first we address the question of how execution is tied to mineness.

Do mineness and execution dissociate? Here it's good to keep clear on the distinction between executing just *any* behaviour and executing an intention in particular. It would seem that one can be aware of executing some behaviour without an accompanying sense of mineness. For instance, in anarchic hand syndrome (AHS), a condition caused by lesion to the supplementary motor area that results in loss of control over the movements of the contralesional limb, the agent is aware of executing various behaviours with that limb, e.g., picking up objects, grasping at things, while lacking the subjective sense that the behaviour is theirs. Such patients say things like, "I know it's me, it just doesn't *feel* like me" (as reported in Marcel 2003: 79). But importantly, this is not a case of executing *willed* behaviour. Typically the anarchic limb does things that are at cross-purposes with the agent's conscious intentions. The agent may be trying to turn the page of a book, for example, while the anarchic hand closes it (Banks et al. 1989: 457). We emphasize that it is the sense of executing *willed* behaviour that we take to capture the executive component of agentic phenomenology.

A natural way of understanding the relationship between mineness and action execution is that awareness of executing one's intention is intimately – perhaps necessarily – connected with mineness phenomenology. This might be because an awareness of intention execution grounds mineness phenomenology in some way (yet to be understood). Or it could be because mineness phenomenology is in part constitutive of the awareness of intention execution.⁶ These questions remain to be settled in future work.⁷

2.2.4 Action Perception

A fourth important aspect of agentic phenomenology involves the primarily perceptual experience of what is happening in the world, with the body, and with certain aspects of the mind, as one acts. Regarding perceptual experience in action, two broad areas are relevant. The first concerns the relative importance of perceptual experience to action control. Following the pioneering work of Milner and Goodale (1995) on conscious vision in action, some philosophers – most prominently Andy Clark (2001) – have suggested that contrary to appearances, conscious perception is of little importance for action control. Clark has maintained that "although it may sometimes seem as if conscious seeing is what continuously and delicately guides our fine-tuned motor activity, such online control may be largely and typically devolved to distinct, nonconscious, visual-input-using systems" (511). Other philosophers have pushed back against this interpretation of Milner and Goodale, arguing instead that while conscious perception is not the only source of information for

⁶ A further question surrounds cases of deviant causation. Could an agent experience both mineness and execution in spite of deviant causal chains linking intention and behavior? In our view, the answer to this question must proceed via examination of the kinds of mechanisms that generate both mineness and execution phenomenology. It is at least possible that these mechanisms do not depend directly on the causal chains linking intention and behavior. If so, one could have the illusion of execution and mineness in spite of deviant causal chains.

⁷ Mylopoulos' (forthcoming) aforementioned treatment of agentic phenomenology can be viewed as an attempt to explain how this might work via agentic thoughts that represent oneself as performing the intended action via the essential indexical, thus inextricably linking the experience of mineness to the experience of execution.

action control, for many action-types it remains functionally important (see Wallhagen 2007, Mole 2009, Shepherd 2016, Briscoe and Schwenkler 2016).

A second area of interest concerns the relationship between the phenomenology of perception and that of action.⁸ In this connection, one interesting question raised recently in Shepherd (2016) is this. What is the relationship between one's proprietarily agentive phenomenology and the perceptual experiences one has while acting? Shepherd argues that in experiences of acting, experiences of trying and perceptual experiences fit together in a certain way.

The experience of acting typically consists of temporally extended experiences from more than one [perceptual] modality. These experiences are easily associated with the action being performed in virtue of the fact that their contents fit coherently into the agent's broader plan for action. And their contents fit coherently in virtue of the fact that they are functionally integrated and structured by what the agent is trying to do. (2016: 436)

On this view, the perceptual experiences that are typically co-conscious with agentive experiences are distinct from, but structured by what the agent is trying to do.

We flag a more radical option, while admitting that it remains unclear how best to explicate it. Consider the much-discussed claim that perceptual experiences are transparent in a certain way. On a standard construal of transparency, it is said that 'We normally "see right through" perceptual states to external objects and do not even notice that we are in perceptual states; the properties we are aware of in perception are attributed to the objects perceived' (Lycan 2015). Might the same thing be true of experiences of acting? On this view, one does not simply experience agentive aspects and perceptual aspects co-consciously. One 'sees through' these elements and experiences to the action itself – the agentive and perceptual aspects are attributed to the same unfolding event, namely, the action. The idea needs further development, but it is worth noting that on such a view, it could be that perceptual phenomenology plays a role in at least some experiences of execution and mineness. Execution, recall, is the experience of actually doing what it is one intends to do. And certainly what one intends to do often constitutively involves events in the body and the world beyond it. On such a view, then, phenomenology associated with mineness, execution, and perception in action would turn out to be very closely related. Such a prospect is tantalizing, but developing it is well beyond the present scope.

2.2.5 Action Assessment

⁸ Here we note two interesting issues that are beyond the present scope. The first concerns attempts to characterize actional aspects of perceptual experience. For example, Siegel (2014) argues that some perceptual experiences contain action mandates – something like commands that the agent do such-and-such. And Nanay (2013) has argued that perceptual experiences attribute action-properties (e.g., attribute climb-ability to a tree) that enable perceptual experiences to play the role of the immediate psychological precursors to action. The second concerns the proper role of attention in agentive phenomenology. One plausible view is that attention plays signature roles in structuring an agent's perception in action (Wu 2011). But it is possible that attention has a broader role than this. We leave this possibility for future research.

A fifth cluster concerns experiences that seem to assess the action or some of its sub-components. Here we have in mind experiences of error in action, experiences of success in action, of control over the action, of the ease (or difficulty) of the action, experiences of fatigue, frustration or anxiety associated with the action, and experiences of engagement in the action, of exhilaration in acting, and more.

The experience-types mentioned above may not form a unified cluster. Consider the experience of error in action. Above our amateur skier recalled sensing, 'with alarm,' that he was sliding backwards. This was not supposed to happen – it signaled a mistake. This experience seems clearly to have a function of assessment, although it is not clear how exactly to parse it – is the experience emotional, cognitive, or something else? Other experience-types in this cluster seem to involve sensory and bodily perception, in combination perhaps with the emotions. Notice the combination of bodily perception with a positive emotion of exhilaration in Hemingway's description of skiing: 'The rush and the sudden swoop as he dropped down a steep undulation in the mountainside plucked Nick's mind out and left him only the wonderful flying, dropping sensation in his body' (1995: 121).

We cannot offer a full account of this cluster here. We do, however, wish to highlight an interesting account of a subset of action assessment experiences, due to Joelle Proust (2013). According to Proust, a certain class of nonconceptual metacognitive experiences plays an important role for mental acts, and by extension for agency (insofar as mental acts are crucial for exercises of agency more generally; see section 1.4 for discussion of metacognition and the sense of bodily agency). Proust calls these 'noetic feelings' – feelings associated with the feasibility of performing a mental act like remembering some content. For Proust these feelings are nonconceptual in the sense that one need not deploy or possess concepts to have them. And these feelings are metacognitive in the sense that the psychological mechanisms that produce them 'have the function of assessing one's cognitive dispositions as such' with respect to their ability to fulfill various epistemic functions (165). That is, these feelings track "cognitive adequacy," that is, the correct evaluation of the resources available/needed for a given mental task, of such and such import' (224). These feelings are critical for initiating, sustaining, and guiding the operation of various mental actions.

In a mental action . . . the subject needs to appreciate the normative status of the output of the mental act: is the name retrieved correct? Has the list been exhaustively reproduced? Here . . . a subject is sensitive to the norms involved in self-evaluation through a global impression, including feelings of fluency, coherence, and so on . . . (155)

There is much more to Proust's account than this. Here, it suffices to note that she has identified an interesting class of experiences that are, if she is right, crucial for the kinds of mental actions that make up a large part of agentive phenomenology. Reflection on Proust's account, and on these kinds of 'noetic feelings' generally, seems like a worthwhile endeavor for the agentive phenomenologist.

2.2.6 Freedom

A sixth cluster revolves around the phenomenology of free will. Philosophers are divided on a number of questions surrounding this phenomenology. Is there a phenomenology of free will for all actions, or only for certain ones (e.g., decisions)

(see Lau and Mylopoulos 2014 for some doubts on both counts)? Is the phenomenology better understood as a component of another cluster (e.g., of mineness: see Horgan, Tienson and Graham 2003), or is it a separate, primitive element of agentive phenomenology? Is there a phenomenology of free will at all? On these questions, there is no consensus.⁹

Certainly the primary question regarding the phenomenology of free will concerns how it maps onto philosophical positions on the nature of free will. Many philosophers appear to believe that the phenomenology is libertarian in the sense that if the experience were veridical, libertarianism – the view that free will is incompatible with determinism, and that we possess free will – would be true. Here is C.A. Campbell in a representative passage.

Everyone must make the introspective experiment for himself: but I may perhaps venture to report . . . that I cannot help believing that it lies with me here and now, quite absolutely, which of two genuinely open possibilities I adopt. (Campbell 1951: 463)

Other philosophers find nothing in their phenomenology suggesting libertarianism, although different reasons are offered for this position. Horgan (2011) argues that introspection is not powerful enough to tell us whether agentive phenomenology includes libertarian contents. Grünbaum (1952) argues that the phenomenology is straightforwardly compatibilist: ‘this feeling simply discloses that we were able to act in accord with our strongest desire at that time, and that we could indeed have acted otherwise if a different motive had prevailed at the time’ (672).¹⁰

Finally, there are those who adopt an intermediate position, according to which our phenomenology involves *both* compatibilist and libertarian elements. Strawson (2010), for instance, maintains that our unreflective stance towards our desires involves the full awareness that we cannot choose them—that they are “just there” (92). And yet, he urges that we have a compatibilist reaction to this, on which “it remains true that their just-thereness is not seen as posing any sort of threat to our freedom” (93). Still, there remains another aspect to our experience that according to Strawson is “profoundly libertarian” (90) in character, centrally involving a sense of oneself as a “self-determining ‘agent-self’” (93) that is “separate from, and somehow irreducibly over and above, all its particular desires, pro-attitudes, and so on” (93).

How might we resolve this dispute? Suppose that there is at least some phenomenology suggestive of free will attached to experiences of deciding and acting. Even if so, it remains unclear how to further characterize this phenomenology: and it is notable in this connection that libertarians typically rest content with the assertion that the phenomenology is libertarian, rather than providing further description. But a range of descriptive options are

⁹ For a more expansive treatment of the phenomenology of freedom, see the appendix in Kriegel (2015).

¹⁰ In a similar vein, Oisín Deery (2015), drawing on a view of David Chalmers’s regarding the content of color experience, has developed an intermediate view on behalf of the philosophical (but not the phenomenological) compatibilist. On Deery’s view, the phenomenology of free will may contain both libertarian and compatibilist contents. The view is offered to the compatibilist as a way of maintaining that the phenomenology of freedom could be veridical even if determinism is true, thus blocking one libertarian route to the existence of libertarian free will. We refer interested readers to Deery’s paper.

available. Is the relevant phenomenology cognitive or thought-like in nature, attached to anticipatory imagery of action-possibilities, perceptual, proprietarily agentive, or something else? Offering a further characterization of this experience may help us to see whether or how the experience could have a content that rules out deterministic causation.

Here is one example of what we have in mind. In a recent paper Chandra Sripada (2016) argues that elements of the phenomenology of freedom are connected with the activity of deliberation. As Sripada has it, in deliberation we go about constructing sets of options relevant to the decision problem(s) we face. In three different senses, Sripada uses the language of space to characterize features of these sets of options. First, some option sets are narrow or spacious in the sense that the relevant options are few or many. Second, some option sets are narrow or spacious in the sense that the distance between the relevant options – roughly, it seems, the different kinds of futures associated with taking one or another option – is great or small. Third, some option sets are narrow or spacious in the sense that the options one could include within them are many or few.

Building slightly on what Sripada says, it appears that different experience-types are associated with these features of option set construction. We sometimes experience our options as narrow or constricting. We sometimes experience what Sripada calls the spaciousness of our options. We sometimes experience what Sripada calls movement between options. When the distance between options is great, Sripada observes that ‘the gap that separates [options] is massive; it is dizzying to traverse it’ (16). And when we recognize that the options we could include in our set are many, we experience what Sripada calls unboundedness.

Sripada stresses that all of this is consistent with compatibilism. Moving beyond Sripada, one might suggest that these elements of the phenomenology of freedom could be deployed in a deflationary explanation of claims regarding libertarian phenomenology. Could it be that the experience of relative unboundedness in deliberation, or relative spaciousness and movement at the time of a decision, are wrongly interpreted as a libertarian ability to do otherwise? The suggestion has at least some plausibility, especially for those who find it difficult to see how phenomenology could be suggestive of indeterministic causation.

3 Conclusion

At the beginning of this chapter we glossed agentive phenomenology as the set of those experience-types associated with exercises of agency, and paradigmatically with intentional action. Given the many ways agents act intentionally, one might expect the relevant experience-types to form a broad and diverse class. Our discussion here bears that out. Empirical work on the sources of agentive phenomenology has made some progress in elucidating the kind of psychological mechanisms that undergird it. But as we noted, this work has often suffered from a lack of clarity regarding the target of explanation. One and the same study might measure and probe multiple aspects of phenomenology simultaneously, some proprietarily agentive and others not. In order to both make salient this concern, and make some progress in addressing it, we divided the experience-types that we take to constitute (exhaustively or at least nearly so) agentive phenomenology into six clusters, which we called the phenomenology of purposiveness, mineness, execution, action perception, action assessment, and freedom. Each cluster raises its own interesting questions, and questions remain about the relationships between the clusters. We are hopeful that progress on these questions may serve to help guide future empirical work.

Sources

- Banks, G., Short, P., Martínez, A.J., Latchaw, R., Ratcliff, G. and Boller, F. (1989), 'The Alien Hand Syndrome: Clinical and Postmortem Findings', in *Archives of Neurology* 46/4: 456-459.
- Bayne, T. (2008), 'The Phenomenology of Agency', in *Philosophy Compass* 3/1: 182-202.
- Bayne, T. (2011a), 'Agentive Experiences as Pushmi-pullyu Representations', in A. A. Aguilar, K. Frankish (eds.), *New Waves in Philosophy of Action* (New York, NY: Palgrave Macmillan), 219-236.
- Bayne, T. (2011b), 'The Sense of Agency', in F. Macpherson (ed.), *The Senses* (Oxford: Oxford University Press), 490-524.
- Bayne, T. (2011c), 'Libet and The Case for Free Will Scepticism', in R. Swinburne (ed.), *Free Will and Modern Science* (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*, Cambridge, MA: MIT Press), 49-68.
- Bayne, T. and Montague, M. (2012), *Cognitive Phenomenology* (Oxford, UK: Oxford University Press).
- Bayne, T. and Pacherie, E. (2007), 'Narrators and Comparators: The Architecture of Agentive Self-awareness', in *Synthese* 159: 475-491.
- Bear, A., and Bloom, P. (2016), 'A Simple Task Uncovers a Postdictive Illusion of Choice', in *Psychological Science* 1: 9.
- Bermúdez, J. L. (2010), 'Action and Awareness of Agency: Comments on Christopher Frith', in *Pragmatics and Cognition* 18/3: 584-596.
- Biran, I., Giovannetti, T., Buxbaum, L., and Chatterjee, A. (2006), 'The Alien Hand Syndrome: What Makes The Alien Hand Alien?' in *Cognitive Neuropsychology* 23/4: 563-582.
- Blakemore, S. J., and Frith, C. (2003), 'Self-awareness and Action', in *Current Opinion in Neurobiology* 13/2: 219-224.
- Blakemore, S.J., Frith, C.D. and Wolpert, D.M. (1999), 'Spatio-temporal Prediction Modulates the Perception of Self-produced Stimuli', in *Journal of Cognitive Neuroscience* 11/5: 551-559.
- Blakemore, S.J., Oakley, D.A. and Frith, C.D. (2003), 'Delusions of Alien Control in The Normal Brain', in *Neuropsychologia* 41/8: 1058-1067.
- Brentano, F. (1874), *Psychologie vom empirischen Standpunkte* (Vol. 1) (Duncker & Humblot).
- Briscoe, R., and Schwenkler, J. (2015), 'Conscious Vision in Action', in *Cognitive Science* 39/7: 1435-1467.
- Campbell, C. A. (1951). 'Is Free Will a Pseudo-Problem?' in *Mind* 60/240: 441-465.
- Carruthers, G. (2012), 'The Case for The Comparator Model as An Explanation of The Sense of Agency and its Breakdowns', in *Consciousness and Cognition* 21: 30-45.
- Carruthers, G. (forthcoming), 'A metacognitive model of the feeling of agency over bodily actions' in *Psychology of Consciousness: Theory, Research, and Practice*.
- Carruthers, P. (2007), 'The Illusion of Conscious Will', in *Synthese* 159/2: 197-213.
- Chambon, V., Filevich, E., and Haggard, P. (2014), 'What is the Human Sense of Agency, and is it Metacognitive?' In *The Cognitive Neuroscience of Metacognition* (Springer Berlin Heidelberg), 321-342.
- Christensen, W., Sutton, J. and McIlwain, D.J. (2016), 'Cognition in Skilled Action: Meshed

- Control and The Varieties of Skill Experience', in *Mind & Language* 31/1: 37-66.
- Clark, A. (2001), 'Visual Experience and Motor Action: Are the Bonds Too Tight?' in *Philosophical Review* 110: 495–519.
- Daprati, E., N. Franck, N. Georgieff, J. Proust, E. Pacherie, J. Dalery, and Jeannerod, M. (1997), 'Looking for the Agent: An Investigation into Consciousness of Action and Self-Consciousness in Schizophrenic Patients', in *Cognition* 65: 71–86.
- de Biran, F.M.M. (1932), 'Essai sur les fondements de la psychologie (1812)', *Edition Tisserand des Œuvres de Maine de Biran, tomes VIII et IX, Alcan.*
- Deery, O. (2015), 'The Fall from Eden: Why Libertarianism Isn't Justified by Experience', in *Australasian Journal of Philosophy* 93/2: 319-334.
- Dow, J. M. (forthcoming), 'Just Doing What I Do: On the Awareness of Fluent Agency', in *Phenomenology and the Cognitive Sciences*.
- Farrer, C., Franck, N., Georgieff, N., Frith, C.D., Decety, J. and Jeannerod, M. (2003), 'Modulating the Experience of Agency: A Positron Emission Tomography Study', in *Neuroimage* 18/2: 324-333.
- Fourneret, P. and Jeannerod, M. (1998), 'Limited Conscious Monitoring of Motor Performance in Normal Subjects', in *Neuropsychologia* 36/11: 1133-1140.
- Ginet, C. (1990), *On Action* (Cambridge University Press).
- Grünbaum, A. (1952), 'Causality and the Science of Human Behavior', in *American Scientist* 40/4: 665-76.
- Haggard, P. and Chambon, V. (2012), 'Sense of Agency' in *Current Biology* 22/10: R390-R392.
- Haggard, P., Clark, S., Kalogeras, J. (2002), 'Voluntary Action and Conscious Awareness', in *Nature Neuroscience* 5/4: 382–385.
- Hemingway, E. (1995), 'Cross Country Snow', in J. Fenton (ed.), *Ernest Hemingway: The Collected Stories*, (London: Everyman's Library).
- Horgan, T., Tienson, J., and George, G. (2003), 'The Phenomenology of First-person Agency', in S. Walter and H. Heinz-Dieter (eds.), *Physicalism and Mental Causation*, (Imprint Academic), 323-340.
- Horgan, T. (2007), 'Agentive Phenomenal Intentionality and the Limits of Introspection', in *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* (New York: Oxford University Press), 57-78.
- Hume, D. (2000), *A Treatise of Human Nature*, D.F. Norton and M.J. Norton (eds.), (Oxford: Oxford University Press).
- James, W. (1880), *The Feeling of Effort* (Vol. 14), (The Society).
- Jeannerod, M. (2007), 'Being Oneself', in *Journal of Physiology-Paris* 101/4: 161-168.
- Kriegel, U. (2015). *The Varieties of Consciousness* (New York: Oxford University Press).
- Lau, H. C., Rogers, R. D., Haggard, P., and Passingham, R. E. (2004), 'Attention to Intention', in *Science* 303/5661: 1208-1210.

- 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', in *Brain* 106: 623-642.
- Lycan, W. (2015), 'Representational Theories of Consciousness', in E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2015 Edition), <<http://plato.stanford.edu/archives/sum2015/entries/consciousness-representational/>>.
- Mandik, P. (2010), 'Control Consciousness', in *Topics in Cognitive Science* 2: 643-657.
- 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* (Oxford: Oxford University Press), 48-93.
- Mele, A. (2009), *Effective Intentions: The Power of Conscious Will* (Oxford University Press).
- Mele, A. (2010), 'Conscious Deciding and the Science of Free Will', in R. Baumeister, A. Mele, and K. Vohs (eds.), *Free Will and Consciousness: How Might They Work?* (Oxford University Press).
- Merleau-Ponty, M. (1996), *Phenomenology of Perception*, C. Smith (trans.), (Motilal Banarsidass Publishe).
- Metcalf, J., Eich, T. S., and Miele, D. B. (2013), 'Metacognition of Agency: Proximal Action and Distal Outcome', in *Experimental Brain Research* 229/3: 485-496.
- Milner, R.D. and Goodale, M.A. (1995), *The Visual Brain in Action* (Oxford: Oxford University Press).
- Mole, C. (2009), 'Illusions, Demonstratives and the Zombie Action Hypothesis', in *Mind* 118/472: 995-1011.
- Mylopoulos, M. (2012), 'Evaluating the Case for the Low-level Approach to Agentic Awareness', in *Philosophical Topics* 40/2: 103-127.
- Mylopoulos, M. I. (2015), 'Agentic Awareness is Not Sensory Awareness', in *Philosophical Studies* 172/3: 761-780.
- Mylopoulos, M.I. and Lau, H. (2014), 'Naturalizing Free Will', in A. Mele (ed.), *Surrounding Free Will: Philosophy, Psychology, Neuroscience* (Oxford University Press).
- Mylopoulos, M. (forthcoming), 'A Cognitive Account of Agentic Awareness', in *Mind & Language*.
- Nahmias, E. (2005), 'Agency, Authorship, and Illusion', in *Consciousness and Cognition* 14: 771-785.
- Nanay, B. (2013), *Between Perception and Action* (Oxford: Oxford University Press).
- O'Shaughnessy, B. (2003), *Consciousness and the World* (Oxford: Clarendon Press).
- Pacherie, E. (2008), 'The Phenomenology of Action: A Conceptual Framework', in *Cognition* 107: 179-217.
- Pacherie, E. (2012), 'The Phenomenology of Joint Action: Self-Agency vs. Joint-Agency', in Axel Seemann (ed.), *Joint Attention: New Developments* (MIT Press).
- Prinz, J. J. (2007), 'All Consciousness is Perceptual', in B. P. McLaughlin and J. Cohen (eds.), *Contemporary Debates in Philosophy of Mind* (Malden, MA: Blackwell Publishing Ltd.).
- Proust, J. (2013), *The Philosophy of Metacognition: Mental Agency and Self-awareness* (Oxford University Press).

Randall, W. (2009), *Another Long Day on the Piste: A Season in the French Alps* (Little, Brown Book Group).

Ricoeur, P. (1966). *Freedom and Nature: The Voluntary and the Involuntary* (Vol. 1) (Northwestern University Press).

Sartre, J. P. (1969), *Being and Nothingness: An Essay on Phenomenological Ontology*, H.E. Barnes (trans.), (New York: Washington Square Press).

Schurger, A., Sitt, J. D. and Dehaene, S. (2012), 'An Accumulator Model for Spontaneous Neural Activity Prior to Self-initiated Movement', in *Proceedings of the National Academy of Science* 109/42: E2904-13.

Schurger, A., Mylopoulos, M. and Rosenthal, D. (2016), 'Neural Antecedents of Spontaneous Voluntary Movement: A New Perspective', in *Trends in Cognitive Sciences* 20/2: 77-79.

Shepherd, J. (2013), 'The Apparent Illusion of Conscious Deciding', in *Philosophical Explorations* 16/1: 18–30.

Shepherd, J. (2015a), 'Scientific Challenges to Free Will and Moral Responsibility', in *Philosophy Compass* 10/3: 197-207.

Shepherd, J. (2015b), 'Conscious Control Over Action', in *Mind & Language* 30/3: 320-344.

Shepherd, J. (2016), 'Conscious Action/Zombie Action', in *Noûs* 50/2: 419-444.

Siegel, S. (2014), 'Affordances and the Content of Perception', in B. Brogaard (ed.), *Does Perception Have Content?* (Oxford: Oxford University Press), 51–75.

Sripada, C. (2016), 'Free Will and the Construction of Options', in *Philosophical Studies*, DOI: 10.1007/s11098-016-0643-1.

Strawson, G. (2010), *Freedom and Belief (Revised Edition)*, (Oxford: Oxford University Press).

Synofzik, M., Thier, P., Leube, D. T., Schlotterbeck, P. and Lindner, A. (2010), 'Misattributions of Agency in Schizophrenia are Based on Imprecise Predictions about the Sensory Consequences of One's Actions', in *Brain* 133/1: 262–71.

Synofzik, M., Vosgerau, G., and Newen, A. (2008), 'Beyond the Comparator Model: A Multifactorial Two-step Account of Agency', in *Consciousness and Cognition* 17/1: 219-239.

Wallhagen, M. (2007), 'Consciousness and Action: Does Cognitive Science Support (Mild) Epiphenomenalism?' in *British Journal for the Philosophy of Science* 58/3: 539–561.

Wegner, D.M. (2002), *The Illusion of Conscious Will* (Cambridge, MA: Bradford Books).

Wegner, D.M. (2004), 'Précis of the Illusion of Conscious Will', in *Behavioral and Brain Sciences* 27/5: 649-659.

Wegner, D. M. and Wheatley, T.P. (1999), 'Apparent Mental Causation: Sources of the Experience of Will', in *American Psychologist* 54: 480-492.

Wenke, D., Fleming, S. M., and Haggard, P. (2010), 'Subliminal Priming of Actions Influences Sense of Control over Effects of Action', in *Cognition* 115/1: 26-38.

Wood, J. M. (2014), 'Hume and the Phenomenology of Agency', in *Canadian Journal of Philosophy* 44/3-4: 496-517.

Wu, W. (2011), 'What is Conscious Attention?' in *Philosophy and Phenomenological Research* 82/1: 93-120.