

# ***TIME TO ACT: THE DYNAMICS OF AGENTIVE EXPERIENCES***

**Elisabeth Pacherie**

Institut Jean Nicod  
ENS, EHESS, CNRS, Paris  
pacherie@ens.fr

*This paper is dedicated to the memory of Marc Jeannerod.*

## **1. Introduction**

Even the briefest of actions, such as pressing a button, flexing the wrist or raising an arm, takes time. Before one even acts, one has to decide what to do and when. When the time to act has come – which may be immediately following the decision, as when one acts on the spur of the moment –, and provided one recognizes that it has come, the action starts and then progresses until it either reaches successful completion, or, as the case may be, is aborted or ends in failure. Obviously then, actions unfold in time, but so do experiences of agency or, as I shall call them, agentive experiences.

Despite the recent surge of interest in the sense of agency among both philosophers and cognitive scientists, the import of the fact that agentive experiences unfold in time remains to this day largely under-appreciated. This is not to say that the temporal characteristics of agentive experiences have been completely neglected. Indeed, one of the main findings of experimental investigations into the sense of agency is the phenomenon of intentional binding, first reported ten years ago by Patrick Haggard and his colleagues (Haggard, Clark & Kalogeras, 2002), whereby a voluntary action and its external sensory consequences are compressed together in subjective time. The effect of voluntary action on subjective time perception isn't however what I'll be primarily interested in. Rather my main concern will be with how agentive experiences evolve in time. In a recent paper, the philosopher John McDowell defends a view of intentions "as a kind of continuant whose instances change their shape as time passes" (McDowell, 2011: 16). This image of a continuant is also, in my view, very helpful when applied

to agentive experiences. Experiences of agency change their shape as time passes, shifting from experiences of being about to act, to experiences of initiating action, progressing through it, and finally reaching completion.

In this paper, I argue that taking the full measure of the time dynamics of agentive experiences can help us better appreciate what their intentional structure is, in what sense agentive experiences are experiences of activity as such, and what form their contribution to self-knowledge and knowledge of action takes. In recent years, a number of empirical models of how and where in the cognitive architecture the sense of agency is generated have been proposed. I also argue that taking into account the dynamic nature of agentive experiences allows us to better understand in what sense these models should be seen as complementary rather than as rival. I start with a preliminary characterization of agentive experiences in section 2. I then present a brief survey of recent empirical work on the sense of agency and discuss integrative models of the sense of agency (sections 3 and 4). In section 5, I consider two main respects in which agentive experiences change their shape as time passes. The first kind of shape-shifting has to do with the contents of these experiences and their level of grain. The second kind of shape-shifting concerns their intentional structure, with a shift in direction of fit. In the final section, I examine in what sense agentive experiences can be sources of knowledge about our own agency and what kind of knowledge they can yield.

## ***2. Agentive experiences***

After a long period of neglect, the experience of agency has finally gained some prominence in the agenda of philosophers and scientists alike. The explosion of interest in the topic in the last decade reveals its complexity. The burgeoning literature on the content and sources of agentive experiences highlights their many facets. A non-exhaustive list of proposed distinctions includes experience of deliberation, experience of decision, awareness of a goal, awareness of an intention to act, awareness of initiation of action, awareness of movement, sense of activity, sense of mental effort, sense of physical effort, sense of control, experience of authorship, experience of intentionality, experience of purposiveness, experience of freedom, and experience of mental causation. To avoid becoming overwhelmed by this terminological profusion, it may be useful, following Bayne (2010), to distinguish between core and non-core elements of agentive experiences. Core elements are elements that must be possessed by any agentive experiences whatsoever and that are essential to agentive experiences, and non-core elements are elements that may but need not be present within experiences of agency.

At the very minimum, to have an agentive experience is to be aware of oneself as acting, that is, to have a sense of agency. It is difficult, however, to conceive of a sense of agency or a feeling

of doing as existing independently of any experience of what it is one is doing, as a feeling simply floating in the air so to speak. It seems that an agentic experience also minimally involves some sense of what it is one is doing, whether it be an experience of moving one's body in a certain way (e.g., moving a finger), of performing an action of particular kind (e.g., pressing a key) or of acting with a certain goal (e.g., starting a new paragraph). Having an agentic experience does not require one to identify the action the experience is about at a particular level of specification, but it demands that the action be specified in some way, however vague the specification is. So I take it that agentic experiences include two core elements: a sense of agency and some specification of the action the sense of agency is for.<sup>1</sup>

Non-core elements of agentic experiences include experiences of effort, experience of deliberation, experience of decision, and other elements that need not always be present within an experience of agency. Arguably, not all actions are effortful or felt as such. Indeed, one finds reports in the clinical literature of people who after brain damage have lost the capacity to experience actions as effortful but still enjoy agentic experiences (Naccache et al., 2005). It is also not uncommon for one to act on the spur of the moment, or for actions, especially routine actions, to be triggered more or less automatically by cues in the environment. We can have agentic experiences for these actions, but these won't include experiences of deliberation or of decision. This is often the case for routine actions that can be triggered more or less automatically. For instance, I found myself a few minutes ago cleaning my glasses. As soon as I became aware of what I was doing, I had an agentic experience for the action I was performing but it did not include among its elements an experience of deliberating whether or not to clean my glasses or even deciding to do so.

For the sake of completeness, one may want to add a third category to these two categories of core and non-core elements of agentic experiences. This category would not be a third category of elements of agentic experiences, but would comprise instead "theoretical" or "philosophical" interpretations of agentic experiences. Arguably, the experience of conscious mental causation and the experience of freedom (at least when interpreted in a metaphysically loaded way) may be regarded as falling in this third category (Pacherie, 2010).

Since I take it that agentic experiences include as a core element some specification of the action for which one experiences a sense of agency, I am obviously committed to agentic experiences having both a phenomenal character and an intentional content.<sup>2</sup> To lay my cards on

---

<sup>1</sup> Bayne (2010) leaves it open whether the agentic core should take the richer form I suggest or can reduce to a bare experience of oneself as acting.

<sup>2</sup> Note that this shouldn't be taken to imply that I do myself take the experience of agency narrowly construed, i.e., as a bare feeling of agency, to be devoid of intentional content. All I am suggesting is that

the table, I should also add that I take the relation between intentional and phenomenal properties of agentic experiences to be sufficiently strong that we can go a long way, if not perhaps all the way, into capturing differences between agentic experiences by appealing to differences in their intentional contents. Admittedly, the phenomenology of agency is rather elusive and the most effective method to get some grip on it may be to search for it under the lamp-post of intentional content. Finally, I happen to believe that the intentional content of agentic experiences is at least in part (and perhaps in a large part) non-conceptual (Pacherie, 2011). However, my commitments stop there and in this paper I take no stand on whether or not phenomenal properties supervene on, or reduce to, intentional properties.

### ***3. The sources of agentic experiences***

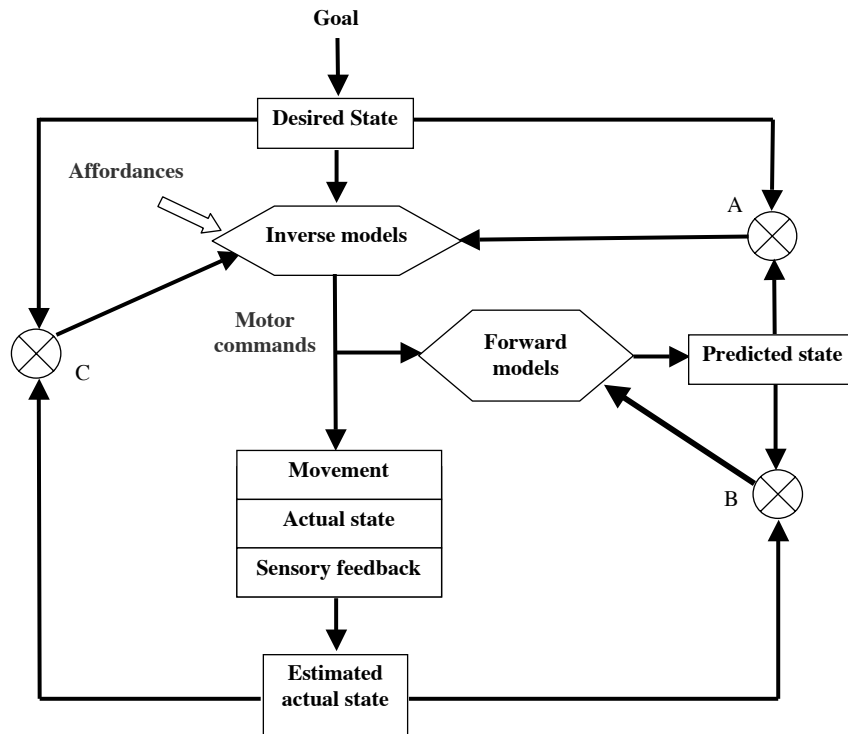
Empirical research on agency has explored a variety of potential cues to agency, and a number of different models of how agentic experiences are generated have been proposed. What these models have in common is, first, a stronger focus on the sense of agency narrowly conceived than on how actions are represented in agentic experiences, and, second, their reliance on a congruence principle: the sense of agency is produced when there is a match between cues  $x$  and  $y$ . The points on which they differ concern (1) the nature of the cues being compared, (2) the nature of the processes involved in the production of the sense of agency, and (3) how closely these processes are related to action production and control processes.

Two theoretical positions define the two ends of the spectrum of possibilities: the motor prediction view and the cognitive reconstruction view. On the motor prediction view, the sense of agency is generated by processes dedicated to action control. On the cognitive reconstruction view, the sense of agency is generated by a general-purpose process of retrospective causal inference.

The motor prediction view is inspired by computational theories of motor control. According to these theories, motor control strategies are based on the coupling of two types of internal models: inverse models and forward models (see Figure 1). An inverse model computes the motor commands needed for achieving a desired state given the current state of the system and of the environment. An efference copy of these commands is fed to forward models, of which there are two classes. The forward dynamic model captures the dynamics of bodily movement and predicts the next state of the system, while the forward sensory model makes predictions about the sensory consequences of the movement.

---

even if one were to argue that the feeling of acting is some kind of raw phenomenal feel, one couldn't argue that this holds for an agentic experience taken as a whole, insofar as this experience must also contain as a core element some sense of what it is one is doing.



**Figure 1. The basic components of a motor control system based on internal models.**  
Adapted from Frith, Blakemore, & Wolpert, 2000.

The control of action depends in a large part on the coupling of inverse and forward models through a series of comparators: error signals arising from the comparison of desired, predicted, and actual states are used for various kinds of regulation. A first kind of comparator (labeled A in Figure 1) takes as input representations of the desired state and of the predicted state and sends an error signal to the inverse model if a difference is found. Such a mechanism can be used to maintain accurate performance in the presence of feedback delays. It can also be used for mental practice and planning, as forward models can predict the sensory outcome of an action without the action being actually carried out. A second kind of comparator mechanism (labeled B in Figure 1) compares the predicted consequences of a motor command with its actual consequences. The result of this comparison can be used to update the forward model and improve its functioning. It can also be used to filter sensory information and to distinguish the component that is due to self-movement from that due to changes in the world (Blakemore et al., 1999). Finally, a third kind of comparison is between desired state and actual feedback (labeled C in Figure 1). Errors derived from the difference between the desired state and the actual state can be used to update the inverse models and improve performance. This kind of comparison is therefore important for motor learning.

According to the comparator model of the sense of agency (Blakemore & Frith, 2003; Frith, Blakemore, & Wolpert, 2000a, 2000b), the comparison of sensory prediction with both desired states and sensory reafferences is also the basis of the sense of agency. The proponents of this model claim (1) that awareness of initiating an action is based on a representation of the predicted consequences of making that action, rather than its actual consequences, and on the congruence between the predicted state and the desired state, and (2) that for this experience of agency to continue, the predicted consequences would also have to remain congruent with the sensory reafferences when they become available.

Claim (1) – and therefore the possibility that the sense of agency can emerge in advance of actual sensory effect and be based on premotor processes alone – is supported by evidence that awareness of initiating a movement in healthy subjects is reported by the agent between 80-200 milliseconds before the movement actually occurs (Libet *et al.*, 1983; Libet, 1985). In an experiment extending Libet's work, Haggard & Eimer (1999) confirmed that movement judgments corresponding to the awareness of movement onset preceded actual movement and also showed that these judgments were unrelated to the general readiness potential but covaried with the lateralized readiness potential, suggesting that awareness of action is tied to the selection of a specific motor program. In another experiment using Transcranial Magnetic Stimulation (TMS), Haggard and Magno (1999) showed that applying TMS over the primary motor cortex created a large delay of the actual reaction time (movement onset) but a much smaller delay of the time of awareness of movement, whereas applying TMS over pre-motor areas, specifically the SMA, led to a much smaller delay of actual reaction time but to a greater delay in the awareness of movement. These data support the view that awareness of movement onset is generated upstream of the primary motor cortex and associated with pre-motor processes. Interestingly, in another study, Sirigu and colleagues (Sirigu *et al.*, 2004) showed that patients with parietal damage could report when they started moving but not when they first became aware of their intention to move. More recently, Desmurget and colleagues (Desmurget *et al.*, 2009) applied electrical stimulation over the parietal cortex in patients undergoing awake surgery and found that it provoked a strong intention to move the contralateral hand, arm, or foot, and when stimulation intensity an illusory experience of actually moving. On the basis of all these results, Desmurget & Sirigu (2009) proposed that a network organized around three major regions – the posterior parietal cortex, the supplementary motor area and the premotor cortex – underlies movement intentions and motor awareness. They therefore argue that motor awareness does not emerge from the sensory signals generated by the movement, but from the predictions we make about the movement before action onset.

Evidence for claim (2) – that the sense of agency also depends on the congruence between predictions and sensory reafferences – comes from studies where these reafferences are artificially manipulated by introducing temporal delays and spatial distortions of feedback. These studies indicate that discrepancies between predictions and sensory reafferences affect tactile sensations (Blakemore, Wolpert, & Frith, 1998; Blakemore, Wolpert, & Frith, 2000), visual perception of one's own actions (Leube et al., 2003), and demonstrate that the sense of agency is gradually reduced as these discrepancies increase (Fournieret & Jeannerod, 1998; Knoblich & Kircher, 2004; Knoblich, Stottmeister, & Kircher, 2004; Leube et al., 2003; Sato & Yasuda, 2005; van den Bos & Jeannerod, 2002).

The results of some of these studies also suggest that the motor system generates predictions at several levels of grain and that we need to distinguish between low-level sensorimotor predictions and reafferences and higher-level perceptual predictions and feedback (Gallagher, 2007; Jeannerod, 2009; Knoblich & Repp, 2009; Pacherie, 2008). The vast majority of our actions aim at producing effects in the environment and we normally attend to the perceptual effects of our movements rather than to low-level sensory reafferences. It may therefore be that perceptual cues rather than sensorimotor cues are crucial to the sense of agency. Direct evidence for this view comes from an experiment of Fournieret and Jeannerod (1998) in which subjects were instructed to move a stylus on a graphic tablet along a straight line to a visual target. Subjects could not see their drawing hand, but its trajectory was visible as a line on a computer screen. On some trials, the experimenter introduced a directional bias electronically so that the visible trajectory no longer corresponded to that of the hand. When the bias was small ( $< 10^\circ$ ) subjects made automatic adjustments of their hand movements to reach the target but remained unaware that they were making these corrections. It is only with larger biases that subjects became aware of a discrepancy and began to use conscious monitoring of their hand movement to correct for it and to reach the target. These results suggest that although discrepancies between predicted and actual sensory feedback are detected at some level since they are used to make appropriate corrections of the hand movement, they are not normally consciously monitored. Rather, the sense of agency for the action seems to rely mostly on a comparison of the predicted and actual perceptual consequences of the action.

These considerations do not call into question the predictive view of the sense of agency. What they suggest is that a more complex, multi-layered view of action control mechanisms is needed, a view that allows for different levels of action-specification and, correspondingly, different levels of prediction of the consequences of an action. The lower level and more fine-grained sensorimotor predictions would be used for automatic motor control and the less fine-

grained perceptual predictions would be used in the conscious control of action and would contribute more to the sense of agency.

In contrast, the cognitive reconstruction view downplays the contribution of the motor system to the sense of agency and proposes that it is inferred retrospectively from the existence of a match between a prior thought and an observed action. Thus, on Wegner's "theory of apparent mental causation" (Wegner, 2002), a general-purpose causal inference process is at play. If an action is consistent with a prior thought of the agent and other potential causes of the action are not present or salient, a sense of agency for the action will be induced.

There is also empirical evidence that high-level inferential processes play a role in determining the sense of agency for an action. Studies of Wegner and colleagues have demonstrated that cognitive cues can alter the sense of agency for an action independently of changes in sensorimotor and perceptual cues. For instance, in their "I-Spy" study (Wegner & Wheatley, 1999), a participant and a confederate of the experimenter had joint control of a computer mouse that could be moved over any one of a number of pictures on a screen. When participants had been primed with the name of an item on which the mouse landed, they expressed a stronger sense of agency for the action of stopping on that object (when in fact the stop had been forced by the confederate). In another experiment, the "helping hand" study, Wegner and colleagues (Wegner, Sparrow, & Winerman, 2004) also demonstrated that it was not even necessary for a subject to actually move to experience a sense of agency. They had participants watch themselves in a mirror while another person behind them, hidden from view, extended hands forward on each side where participants' hands would normally appear and performed a series of movements. When participants could hear instructions previewing the movements, their sense of agency for these movements was enhanced, but such vicarious agency was not felt when the instructions followed the movements.

Further studies also suggest that subliminally priming an outcome just before the outcome is produced can enhance the sense of agency for that outcome (Aarts, Custers, & Wegner et al., 2005) and that priming an outcome relatively far in advance can augment self-agency as well, but only if the outcome is attached to positive affect (Aarts, Custer, & Marien, 2009). Religious beliefs have also been shown to have an influence on the sense of agency. Dijksterhuis and colleagues (Dijksterhuis et al., 2008) investigated how subliminally primed thoughts of an agent prior to action can affect the sense of agency for that action in a task where authorship is ambiguous. Participants reported a greater sense of agency when primed with first person singular pronouns. In contrast, the sense of agency decreased when participants were primed with "God", but only among believers.

#### **4. Integrative frameworks**

There is now a growing consensus that the motor prediction view and the cognitive reconstruction view are not mutually exclusive but complementary and that intrinsic cues (cues provided by the motor system) and extrinsic cues (such as cognitive primes) both contribute to the sense of agency (Bayne & Pacherie, 2007; Gallagher, 2007; Knoblich & Repp, 2009; Pacherie, 2008; Sato, 2009; Synofzik, Vosgerau, & Newen, 2008).

The new challenge for researchers is therefore to develop integrative frameworks. To do so, they must answer a number of questions. How do all these agency cues interact? What is their relative weight with regard to experiencing agency? Can extrinsic cues overrule perceptual and sensorimotor evidence? To what extent can the relative weight of different agency cues be modulated by the nature of the task, by the attentional state of the agent, or by the agent's level of expertise?

Some recent studies have started probing these issues. In their "space pilot" study, Metcalfe and Greene (2007) had participants play a computerized game in which they tried to touch, with a cursor, downward scrolling Xs and avoid touching Os and where a number of variables affected accuracy. By manipulating these variables (speed, feedback delay, random discrepancy from mouse input), they were able to show that the degree of subjective control reported by their participants was not just a function of their judged performance (extrinsic cues) but was also sensitive to the variables affecting consistency between motor actions and their effects (intrinsic cues).

Haggard's initial interpretation of the intentional binding phenomenon, whereby voluntary actions and their effects are attracted together in subjective time, was strictly predictive (Haggard et al., 2002; Haggard & Clark, 2003). The idea that intentional binding relies on motor prediction processes is supported by the fact the binding effect appears to be specific to intentional action, since when a sensory consequence is the result of an involuntary movement, the binding effect is reversed and the interval between movement and consequence dilated in subjective time. Further support for the predictive interpretation comes from a study by Engbert and Wohlschläger (2007), where the probability that the action would produce an outcome (a tone) was varied. They found that intentional binding was stronger in the higher outcome probability condition (80%) than in the lower outcome probability condition (20%).

In more recent studies, Haggard and his colleagues investigated the possibility that both prediction and retrospective inference play a role in intentional binding. Moore and Haggard (2008) manipulated the probability that an action would yield a certain outcome. Their results

suggest that when the probability of the outcome given the action is high, predictive processes contribute more to intentional binding, but that when this probability is low, retrospective inference is involved in intentional binding. Moore et al. (2009) further demonstrated that increasing action-outcome contingency (defined as the probability of an outcome  $O$  given an action  $A$ , minus the probability of  $O$  in the absence of  $A$ ) significantly increased the contribution of both predictive and retrospective inferential processes to intentional binding. Since contingency is an index of the causal relations between events, this suggests that the human mind builds internal models of action–effect relations and that intentional binding is sensitive to these causal relations.

In another study, Moore, Wegner and Haggard (2009) investigated the effects that priming subjects with prior thoughts about an effect had on intentional binding. Participants either made a voluntary key press movement or had an equivalent involuntary movement applied passively to their finger. All movements were followed by either a high or a low-pitched tone, at random. In addition, one of these two tones could be presented as a prime, prior to the movement. That prime was either congruent or incongruent with the effect of the ensuing action. Participants were asked to give verbal estimates of the duration of the interval between a key press movement and a tone. In line with previous research, intervals initiated by voluntary movements were perceived as shorter than those involving passive involuntary movements. Prime congruence also modulated the sense of agency for both voluntary and involuntary movements, with congruent prime–effect pairings leading to an enhanced sense of agency (lower interval estimates). However, this modulation by prime congruence was significantly stronger for involuntary than for voluntary movements. These results corroborate the idea that the sense of agency depends on a mechanism that uses multiple cues, both intrinsic and extrinsic to the motor system, as to the origin of actions and their effects. They also suggest that extrinsic cues (e.g. conscious thoughts evoked by primes) are most effective when intrinsic cues (motor signals) are weak or absent.<sup>3</sup>

Sato (2009) investigated the respective contribution to the sense of agency of conceptual congruence between preview information and actual outcome and of sensorimotor congruence between prediction and actual sensory feedback. He found that both of these factors contributed to the sense of agency, but that sensorimotor congruence appeared to have a more robust impact. Knoblich and Repp (2009) investigated how skill level and task difficulty can affect the relative weight that is given to different types of agency cues. Their results indicated that subjects appear to rely more on perceptual cues if a task is difficult or unfamiliar, but that they

---

<sup>3</sup> For a review of work on intentional binding and the sense of agency, see Moore & Obhi (in press).

become more sensitive to sensorimotor cues when the task is less challenging or their expertise with it greater.

Two main approaches have been proposed for building an integrative framework. One is the multifactorial weighting model of the mechanisms underlying the sense of agency initially proposed by Synofzik, Vosgerau, and Newen's (2008). The other is a generalized, hierarchical comparator model (Fletcher & Frith, 2009; Pacherie, 2008). The main difference between these two approaches is that the latter explicitly retains the assumption that predictive processes play an essential role in the production of the sense of agency and therefore that the processes that generate agentic experiences are closely related to the processes involved in action preparation and control. Both models have been criticized for being under-constrained. For instance, Carruthers (In Press) argues that the multifactorial weighting model is virtually unfalsifiable, as it seems that the weights assigned to different factors can always be adjusted, case by case, to explain the data. What this model lacks, according to Carruthers, is some principled way of assigning weights to different agency cues. On the other hand, Wong (In Press) objects to the generalized comparator model that it empties the notion of a comparator of any specific meaning. In particular, he points out that while the comparator mechanism was originally tied to very specific physiological mechanisms, it seems now to be "just standing proxy for whatever mechanism is responsible for the sense of agency insofar as some feedback is involved" (Wong, In Press: 3).

In response to the objection addressed to the multifactorial weighting model, Vosgerau and Synofzik (In Press) acknowledge that the question of how weights are assigned to different cues was not tackled in sufficient detail in their initial presentation of the model, but point out that it can be given a principled answer by appealing to the notion of Bayesian optimal cue integration. The idea is that the sense of agency is determined by an optimal combination of different agency cues, according to their reliability (Hendricks, Wiggers, Jonker, & Haselager, 2007; Lau, Rogers, & Passingham, 2007; Moore, Wegner, and Haggard, 2009; Moore & Fletcher, In Press). In a nutshell, optimal cue integration theory says that to each agency cue is associated a likelihood function that assign it a weight proportional to its reliability and that cue integration is itself the product of the likelihood functions associated to each of the available cues.<sup>4</sup>

It is important to note that the multifactorial weighting model in its new version is not the only model to rely on Bayesian cue integration. Bayesian cue integration is also part and parcel of the generalized comparator model. With respect to Wong's objection, while one can

---

<sup>4</sup> See, for instance, Moore and Fletcher (In Press) for accessible examples of how optimal cue integration works.

acknowledge that in the generalized comparator model, the notion of a comparator ceases to be uniquely tied to very specific physiological mechanisms (although these ties are retained for low level comparators) it is simply untrue that 'comparators' now simply stand proxy for whatever mechanism is responsible for the sense of agency insofar as some feedback is involved. First, as Wong (2012) himself points out, the original comparator model attributed a role to both comparisons of predicted states with actual states (feedback) and comparisons of predicted states with desired states. In the generalized comparator model, the two kinds of comparisons also play a role. It is thus untrue that the generalized comparator model is obsessed with feedback and that comparators stand proxy for any mechanism *provided feedback is involved*. Second, the mechanisms entering the generalized comparator model are constrained in that they involve a hierarchy of prediction and model building. The generalized comparator model can therefore be seen as an instance of a Bayesian hierarchical model. As explained by Fletcher and Frith (2009): " According to this theory, building models of the world through prediction errors is a common feature of many brain systems, and these systems are arranged in a hierarchy so that the prediction error emitted by a lower-level system becomes the input for a higher-level system. At the same time, feedback from the higher-level system provides the prior beliefs for the lower-level system." (2009: 54). Another important feature of the Bayesian hierarchical approach is that, as we move up in the hierarchy, we also move away from predictive processes specifically dedicated to motor control towards more general purpose predictive processes.<sup>5</sup>

What I would like to suggest is that little is gained by opposing a multifactorial weighting model cum optimal cue integration to a generalized comparator model incorporating a Bayesian hierarchical approach. Rather, combining them gives us a better hold on what makes a given agency cue more or less reliable. Vosgerau and Synofzik (In Press) propose that the reliability of an agency cue, hence the weight it is assigned depends on the context. Indeed, it would make little sense to deny that visual feedback, for instance, should be assigned less weight when an action is performed in the dark than when the same action is performed in broad daylight. The further insight provided by the generalized comparator model is that the reliability of an agency cue also depends on structural factors. By structural factors, I mean the level in the Bayesian hierarchy to which an agency cue pertains. As I remarked above, the higher up we move in this Bayesian hierarchy, the looser the ties of prediction processes to action preparation and control processes. It follows that, *ceteris paribus*, low-level agency cues, intrinsic to the motor system, are more reliable indicators of actual agency than high-level cognitive cues. This is reflected, for instance, in the findings of Moore, Wegner and Haggard (2009), described above, that when an

---

<sup>5</sup> It is worth noting that the idea of hierarchical comparator models was already discussed in Jeannerod (1995, 1997).

internal agency cue exists within the motor system alternative external cues to agency have a reduced impact on agentic experience.

Finally, time is one further factor influencing the weight assigned to various agency cues. The cue integration process takes place over time. Different cues may be available at different times during the course of an action and the reliability of a cue at a given time may be a function of which other cues are also available at that time and change when other cues become, or cease to be, available. As we shall now see, this temporal factor comes to the fore when one examines the dynamics of agentic experiences, and the shifts in intentional content and structure characteristic of this dynamics.

### ***5. The dynamics of agentic experience: content and structure shifts***

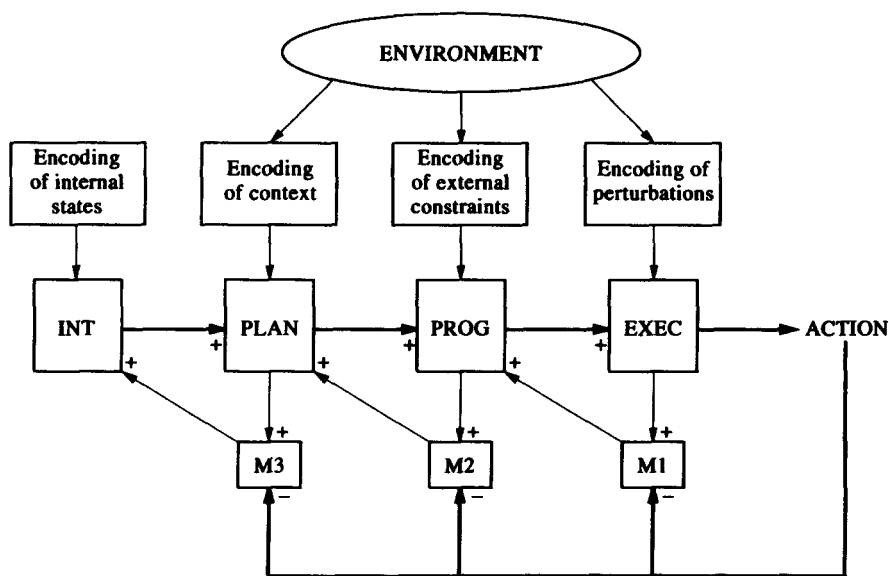
Agentic experiences evolve in time and change their shape as they do. This is what I call the dynamics of agentic experiences. My aim in this section is to characterize this dynamics by focusing on two types of changes agentic experiences undergo: first, their intentional contents change and, second, their intentional structure and direction of fit also changes.

I start with content-shifts. I proposed in section 2 that agentic experiences include two core elements: a sense of agency and some specification of the action the sense of agency is for. I pointed out, however, that having an agentic experience does not require one to identify the action the experience is about at a particular level of specification.<sup>6</sup> Not imposing such a requirement makes sense if indeed the content of an agentic experience shifts as the action it is about unfolds. One aspect of this content-shift involves changes in tense. One can have an experience of being about to *A*, of initiating *A*, of being in the middle of *A*-ing, of nearing the completion of *A*-ing and of having just *A*-ed. If one considers actions somewhat more complex than simply flexing a wrist or raising an arm, it becomes clear that multiple tenses can in principle co-exist within the same agentic experience. Suppose that I am preparing a chocolate mousse, my agentic experience at a given moment may be of having melted the chocolate and being about to whip the egg whites. As the chocolate mousse example can also serve to illustrate, another aspect of content-shift concerns the level of grain at which the action is identified. I can experience my action as preparing a chocolate mousse, whipping the eggs white, rotating my wrist with great energy (which in my case would be a lie, as I am lazy and use an electric hand mixer...), and so on.

---

<sup>6</sup> This, by the way, may be one reason why some people can be skeptical that the core of agentic experiences involves anything more than a bare experience of oneself as acting.

All of this sounds commonsensical enough, or so I hope. But can we say more about what makes one identify an action at a given level of grain, or what makes one concentrate on the past, the present or the future aspects of the action being performed? As I remarked at the outset of the section 3, empirical work on agentic experiences tends to be more concerned with investigating the cues and the processes that generate the sense of agency than with studying how actions are represented in agentic experiences. Nevertheless, it is quite plausible that the various cues that contribute to the sense of agency also contribute to an agent's awareness of what he or she is doing. As the discussion in sections 3 and 4 indicates, a variety of agency cues contribute to the sense of agency but these cues are not all available at the same time. In the early stages of an action, the available cues are mostly prospective cues and as the action progresses, retrospective cues, in the form of feedback and signals arising from the comparison of feedback with predictions, become available in turn. A complementary idea, central to the generalized comparator model discussed in section 4, is that predictions and comparisons of predictions with feedback take place at a number of levels arranged in a hierarchy. This idea was captured by Jeannerod (1995, 1997) in the form of a flow chart model, reproduced as Figure 2.



**Figure 2. A flow-chart model of the ongoing regulation of action.**  
 Reproduced from Jeannerod, 1995.

In his comments on this flowchart, Jeannerod (1995, 1997) makes several important points. First, while the main flow of information goes through the cognitive steps of action specification (e.g., intending, planning, programming and executing), Jeannerod notes that the impression given by this flowchart representation that these steps are organized serially must be corrected. In other words, while the initial component or segment of an action is prepared and executed, the preparation of the next components of the action is already engaged. As he puts it, "the diagram,

at least for what concerns its distal steps, should be considered as one of many parallel lines with a similar structure" (1997: 172). Second, the predictions made are stored in as many short-term memories (represented as M1, M2 and M3 in Figure 2) as there are levels, and reafferences documenting the current state of the action are later fed into the memories. If the action has been completed, the reafferences and the content of the memories will coincide, and the latter will be erased. If the action is incomplete, the memories remain active, which reactivates the corresponding module and generates corrections.

The idea of a set of short-term memories storing predictions and remaining activated until matching reafferences erase their contents gives us a way of understanding which representations of an action are active at a given moment and can thus potentially enter the contents of agentic experiences. This model suggests certain constraints on which action representations can contribute to agentic experience at a given moment, and suggest more generally that our agentic experiences should tend to be present and future oriented, given that short-term memories for actions or action components already completed are promptly erased. Yet, it still implies that many representations of an action can be simultaneously active, first because, as Jeannerod puts it, in a given line, corresponding to the preparation of a given component or segment of an action, several memories storing representations of this segment at different levels of specification can be active, and second, because this line is only one of possibly many such lines where memories are also currently active. Yet, it seems unlikely that all the action representations active at a given time contribute, or at least contribute equally, to agentic experiences.

Several other factors can further influence which action representations enter the contents of agentic experiences. One important factor is attention, which can be voluntary or involuntary. On the one hand, we can voluntarily allocate attention to some aspects of an action more than to others. For instance, when I play a serve in tennis, I can choose to attend to distal aspects of the action, (e.g., where in the court the ball lands), or to more proximal aspects of the action (e.g., how I move my arm and shoulder). On the other hand, stimulus events can also involuntarily capture attention. The short-term memories discussed by Jeannerod are short-lived, since they are erased as soon as matching reafferences become available. The longer they last, the more likely it is that they will capture attention, since their longevity is a sign of trouble. One further factor is the level of skill of the agent. When an action is well rehearsed, the transition between action segments or components can proceed more or less automatically. If, however, the agent is not used to performing actions of that kind, he or she may have to consciously consider what comes next when a given segment is completed. Yet another factor is the agent's preferred level of action identification. According to Action Identification Theory

(Vallacher & Wegner, 1985, 1987, 1989), any behavior can be represented at multiple levels but there are individual differences in agents' preferred level of action identification. Some people show a general tendency to (consciously) represent their behavior at a low level, i.e., in terms of the concrete, mechanistic aspects of their action (e.g. turning a doorknob), whereas other people show a general tendency to represent their behavior at a higher level, i.e. in terms of the goals or purposes of their action (e.g., opening the door or going out). While other factors can also influence the level at which an action is represented, still, *ceteris paribus*, agents whose preferred level of action identification is low will tend to experience their action in terms of producing sensorimotor consequences or outcomes, while agents whose preferred level of action identification is high will tend to experience their action in terms of the goal it serves.<sup>7</sup>

To recap, Jeannerod's hierarchical model of action organization suggests that the action representations that contribute to the contents of agentive experiences are predominantly forward-looking, that is, they are mostly representations of what the agent is currently doing or about to do. It also suggests that active action representations will initially be more abstract and later gain in specificity as one moves from intending to planning, programming and finally executing. As the action nears completion, a converse movement will take place as short-term memories of the predictions associated with execution, programming and planning are successively erased. In addition, while an action progresses, a number of action representations are simultaneously active; which of them enter the contents of agentive experiences depends on a number of factors, including voluntary and involuntary attention, the agent's level of skill and the agent's preferred level of action identification.

Let me now turn to the second type of change agentive experiences undergo. This change can be described as a shift in structure. By structure, I mean the type of relation that holds between these experiential states and states of the world.

Bayne (2008) distinguishes four possible conceptions of the structure of agentive self-experiences. As a first possibility one may think of them as having a thetic structure. On this view, they are supposed to say how things are, have a mind-to-world direction of fit, a world-to-mind direction of causation and veridicality conditions. Alternatively, one may hold that agentive experiences have a telic structure, representing how the world is to be changed and

---

<sup>7</sup> See for instance Van der Weiden *et al.* (2010) for a study of how the effects of priming on agentive experiences crucially depend on action identification level and Belayachi & van Linden (2009) for a discussion of the role of low-level action identification in checking behavior in subjects with obsessive-compulsive disorder.

having a world-to-mind direction of fit and a mind-to-world direction of causation.<sup>8</sup> On the telic account then, these states would have conditions of satisfaction but not veridicality conditions. A third possibility is that agentive experiences are akin to Millikan's *pushmi-pullyu* representations (Millikan, 1995), having both a mind-to-world and a world-to-mind direction of fit and thus being both thetic and telic. A final possibility would be that agentive experiences lack intentional content altogether, involving only raw phenomenal feels. As pointed out by Bayne, although Searle (1983) did much to put the telic account on the table, the descriptive account seems to be the majority view<sup>9</sup>. To the best of my knowledge, only Bayne (Bayne, 2010) and I (Pacherie, 2010, 2011) have tried to explore, and argue for, the third possibility. Finally, while the raw feels approach may perhaps apply to some non-core elements of agentive experiences, it is clearly inapplicable to core agentive experiences, if, as I have argued here, these core experiences essentially involve intentional content. Therefore, I won't discuss it further here.

The view I want to defend here is that the structure of agentive experiences evolves in time, being initially telic, then both telic and thetic (thelic, if you wish), and finally thetic as the action is completed.

I start with a discussion of the constraints states should meet to be considered telic rather than thetic and consider first temporal constraints. No particular temporal constraints on their contents are imposed on states in virtue of their having a thetic structure. For instance, my beliefs can be about past, present or future states of affairs. When temporal constraints apply, they do apply not in virtue of the state having a descriptive structure but in virtue of other features of its mode (as in regret that can only concern past states of affairs). In contrast, there are temporal constraints on the contents of states with a telic structure. I can intend to do something now or to do something tomorrow, but I cannot intend to do something yesterday. These constraints stem from the fact that these states have a mind-to-world direction of causation and from the time-asymmetry of causation: effects cannot precede their causes. For agentive experiences to have a telic structure, it would therefore be necessary, although not sufficient, that they satisfy these temporal requirements. In other words, it would be necessary that these experiences occur before the state of affairs they are about.

In sections 3 and 4, we saw that a variety of processes contribute to the formation of agentive experiences. On the one hand, predictive processes within the motor system yield

---

<sup>8</sup> The labels "thetic" and "telic" were, to my knowledge, first proposed by Humberstone (1992). I adopt his terminology here, as it is much less cumbersome than talk of "mind-to-world" vs. "world-to-mind" direction of fit. For good measure, I also introduce the label "thelic" for states that have both directions of fit.

<sup>9</sup> For recent defenses of the descriptive account, see Bayne (2009) and Proust (2003).

predictions about the consequences or outcome of the prepared action and insofar as agentic experiences are based on these predictions and incorporate their content in advance of actual action they meet the temporal constraints associated to telic states. For instance, experimental evidence indicates that the experience of initiating an action precedes the actual action onset. In contrast, agentic experiences yielded by retrospective inferential processes of the kind proposed by Wegner are post-hoc reconstructions and can only be considered as having thetic structure.

Yet, meeting the temporal constraints just discussed isn't enough for a state to count as telic. If I believe that tomorrow it will be nice and sunny, my belief meets the temporal constraints for telic states, yet is clearly has a thetic structure. My beliefs about the future weather have no influence whatsoever on what the weather will be. For a state to count as telic it is further required that it play a role in bringing it about that the state of affairs it represents obtains. In contrast to weather predictions, the predictions made by the motor control system play an active role in the regulation of ongoing actions. Within the architecture of action control, they have two main regulative roles to play. First, predicted states are compared to desired states. If the comparison yields a match, the gates to action execution are opened; if not, the inverse models are asked to go back to work. Predictions therefore control action initiation. Second, predictions are also compared with reafferences. A match signals completion and the motor system moves on to implement the next step of the action or stops its activity if the whole action has been completed; a mismatch triggers corrections. The role of predictions therefore does not stop at controlling the initiation of actions; rather predictions are involved in the regulation of action during all the time it unfolds. Given their role in the ongoing regulation of actions, predictions can clearly be seen as having telic structure. However, what makes it possible for them to play their role in action regulation is the fact that they also have thetic structure. The forward models that compute these predictions are in the business of extracting from statistical regularities the causal relations between actions and their effects, and using this information to make predictions. In other words, forward models aim at truth; they aim at producing predictions that are accurate. So the predictive states generated within the motor control system should be considered as thetic states, that is states with both thetic and telic structure, whose success in the regulation of actions depends on their correctness.

These considerations suggest the following general picture of the structural shift agentic experiences undergo. Agentic experiences would start their life as possibly telic, at least if one consider that an experience of intending to *A* qualifies as an agentic experience. This, in my view, is problematic if the intention in question is an intention to *A* in the future, but may be defensible for an intention to *A* now. For a large part of their time course, their structural shape

would be largely thetic as they would mostly be based on the predictions made by the motor control system. In the later part of their course, their shape would change to include stronger thetic components as reafferences become available. How strong these thetic components will be will depend on how well reafferences match predictions. When reafferences match predictions, our agentic experiences remain based on predictive information, but when reafferences are unexpected, retrospective processes come into play and transform our agentic experiences to reflect the actual relation between our action and its consequences.

## **6. Knowledge of agency**

Let me close this paper with some thoughts on what agentic experiences can contribute to our knowledge of our own agency. First, to contribute knowledge of any kind, agentic experiences have to be veridical in the first place. I will therefore start with some remarks on the reliability of agentic experiences. Second, one long-standing philosophical debate about self-knowledge concerns the existence of an asymmetry between the knowledge we have of (some aspects of) our mental states and our actions from a first-person perspective and knowledge from a third-person perspective. As pointed out by Eilan and Roessler (2003), with regard to our knowledge of our own agency it is tempting to appeal to agentic experiences to support the claim that there is an asymmetry between agentic knowledge 'from within' and agentic knowledge we may have from a third-person perspective, without falling into the extreme of equating first-person agentic knowledge with introspective knowledge. A third issue concerns whether and how agentic experiences can be experiences of activity as such. As these last two issues are closely related, I will discuss them together.

How reliable are our agentic experiences? If we are to believe Wegner (2002), the answer is that they are largely illusory and shouldn't be taken as reliable guides to our actual agency. Or, at least, this is what Wegner says in his more skeptical moments. There are, however, reasons to be more sanguine than Wegner is. First, we should note that since the issue we are now examining is whether agentic experiences can be a source of knowledge of our own agency, asking whether agentic experiences are reliable amounts to asking whether they are generally accurate or veridical. For this question to make sense, agentic experiences must have structure that is at least in part thetic. If, as we argued in the previous section, agentic experiences have both components with a thetic structure and components with a thetic structure, we have reason to believe that the question indeed makes sense.

Although in actual practice, the distinction may not always be clear-cut, an agentic experience can in principle be inaccurate in two ways: one can mistakenly take oneself to be the agent of an action or one can be mistaken as to the identity of the action one is performing.

Whether and to what extent we can trust an agentive experience to be accurate in either respect depends on whether its content was contributed by predictive or retrospective processes as well as on how rich its intentional content is. To the extent that that an agentive experience is contributed by predictive processes from within the motor control system, we are largely warranted in taking at face value our feeling of doing, since efferent signals trigger these predictive processes. In contrast, no such warrant exists when agentive experiences are the result of retrospective processes where agency is inferred from the existence of a match between perceptual consequences and prior thoughts. This is illustrated by the "helping hands" experiment, devised by Wegner and colleagues (Wegner *et al.*, 2004) and described in section 3, where participants reported a sense of agency for arm movements that were actually performed by another person, provided they had been primed with previews of these movements.

With respect to the identity of the action being performed, information delivered by predictive processes is only as accurate as the predictive models themselves are. As we discussed in section 4, predictive models form a hierarchy. Predictive models at the lower end of the hierarchy would typically be more accurate than predictive models at higher levels in part at least because the causal information they use is extracted from statistical information about motor commands and bodily movements that is more comprehensive and less likely to be biased than the statistical information about actions and their distal effects higher-level predictive models exploit. Thus, suppose I am moving my arm towards a switch. The content of my agentive experience may be that I am turning the light on, if I identify the action at a high level, or that I am extending my arm, if I identify it at a lower level. The former description of my action is more likely to be incorrect (the switch may command the extractor fan rather than the light or the lightbulb may dead), than the latter, as my low-level predictive model normally contains very robust causal information about the relationships between motor commands and bodily movements.

Retrospective inferential processes, in the way they are put to work by Wegner in many of his studies, tend to lead to disaster (i.e., illusory agentive experiences) because all the materials these processes exploit, namely prior thoughts and perceptual information, are extrinsic to the motor control system. There is, however, another way to put these retrospective inferential processes to work that make them contribute to the veracity of agentive experiences rather than undermine it. As proposed by Moore and Haggard (2008), retrospective inferential processes may also kick in when predictions are invalidated by reafferences, in order to ensure that our agentive experience is a correct reflection of the actual relations between our actions and their consequences. What these retrospective inferential processes do, in such a case, is tie efferent

signals to their actual effects in the world. In so doing, they do not just improve the accuracy of our current agentic experiences, they also contribute to improving the accuracy of predictions in the future.

To recap then, the situations in which agentic experiences are likely to be unreliable are, on the one hand, those where they are generated by retrospective processes on the sole basis of cues extrinsic to the motor control system and, on the other hand, those where predictive processes yield incorrect predictions and no sensory reafferences are available to put agentic experiences back on track. Wegner's experiments exploit situations of the former kind and several researchers (Frith *et al.*, 2000; Desmurget & Sirigu, 2009) have argued that patients with anosognosia for hemiplegia with illusory agentic experiences of moving their paralyzed limbs are in the latter kind of situation. However, in ecological situations and provided the predictive systems of the agents are not impaired, predictive and retrospective inferential processes should cooperate to produce veridical agentic experiences.

Let me now turn to a final question: what form does the contribution of (veridical) agentic experiences to our knowledge take? This question echoes a long-standing philosophical debate regarding the existence or not of an asymmetry between knowledge from a first- and from a third-person perspective. As Eilan and Roessler put it:

Traditionally it has been held that we have knowledge of (certain aspects) of our actions 'from within', or from the first-person perspective, on a basis that differs from the kinds of evidence we use in finding out about others' actions. One extreme position would be to reject the very idea of a first-person/third-person asymmetry in relation to knowledge of agency – to assimilate the agent's perspective to that of an outside observer. At the other extreme lies the traditional account of that asymmetry. On this view, knowledge 'from within' is to be equated with introspective knowledge, the only possible objects of which are thought to be 'inner' conscious states of events. Knowledge from within would be knowledge of such things as acts of the will or tryings". (Eilan & Roessler, 2003: 16)

Eilan and Roessler judge, and I with them, that neither of these two extreme positions is very palatable. The question then arises of whether we can steer a middle course, that is, maintain that there is an asymmetry between first- and third-person knowledge of agency and yet avoid construing first-person knowledge of agency as introspective knowledge. Eilan and Roessler also describe Anscombe's account of our first-person knowledge of agency in terms of practical knowledge (Anscombe, 1957) as an attempt to steer such a middle course, but consider, with many other commentators, that her attempted solution is ultimately unsatisfactory. They suggest that we might be more likely to succeed at steering such a middle course by appealing to agentic experiences, although success in this enterprise would still ultimately depend on how agentic experiences themselves are construed.

I agree with Eilan and Roessler that Anscombe's solution is unsatisfactory. It contains, however, an important insight that I would like to highlight before turning one last time to agentive experiences. Anscombe takes agent's knowledge to be a form of practical knowledge. The agent of an action is aware of what she is doing in virtue of rationally controlling her action, where the rational control of an action is matter of practical reasoning leading to a decision to act. Practical reasoning yields practical knowledge and what separates practical knowledge from speculative knowledge is its direction of fit: practical knowledge, in Anscombe's view, is telic rather than thetic. As Moran points out in his discussion of Anscombe's account of practical knowledge, "Unlike theoretical or speculative knowledge, practical knowledge will not be passive or receptive to the facts in question, but is rather a state of the person that plays a role in the constituting of such facts" (Moran, 2004: 47). This captures the important point I wished to highlight in Anscombe's account: insofar as our knowledge of our agency is telic – as practical knowledge indeed is according to Anscombe's construal –, it is knowledge of activity in a strong sense that distinguishes it from third-person knowledge of agency. Where, in my view, Anscombe errs is in taking agentive knowledge to be purely telic, for it seems that this knowledge is then independent of what actually happens. As McDowell points out: "If I express an intention in action by saying 'I am crossing the street', it is not by virtue of something that is so anyway, independently of the intention I express, that I am stating a bit of knowledge, if I am. If it were not for the intention, there would not be the relevant fact. *But even so, if I am not crossing the street what I say is false in the ordinary way; things are not as I say they are.*" (2010: 429; my italics). In other words, my having practical knowledge that I am crossing the street does not guarantee that I am crossing the street, and if I am not, there's a clear sense in which I do not know that I am crossing the street.

This leads me to my final point. Eilan and Roessler propose that we appeal to agentive experiences to steer a middle course in the first-person/third person asymmetry debate, that is, preserve the idea that there is something special about first-person knowledge of agency without construing it as introspective knowledge. For an appeal to agentive experiences to really enable us to steer this middle course, agentive experiences would have to have both thetic and telic structure. If agentive experiences are purely thetic states, then we will not be in a position to capture what makes first-person knowledge of agency special. Here the problem is not, contrary to what Korsgaard seems to think, that "to experience something is (in part) to be passively receptive to it, [that] therefore we cannot have experiences of activity as such" (1996: 204). In response to this objection, Bayne (2008) is quite right to point out that it does not follow from the passivity of experiences that they cannot encode agentive contents. The problem with conceiving of agentive experiences as purely thetic is rather that this cannot ensure that the

asymmetry between first- and third-person agentic knowledge is preserved. If agency is simply a feature of the content of thetic experiential states, then it isn't clear why it couldn't be possible to experience the agency of others in the same way. It certainly seems that we can have an experience of some other agent as actively engaged with the world. If, on the other hand, agentic experiences were purely telic, we would fall prey to the same problem Anscombe's account does. First-person knowledge of agency based on agentic experiences would be bizarrely independent of what actually happens. But then again, if I experience that I am crossing the street but am not actually crossing the street, I don't *know* that I am crossing the street.

If, as I have argued, agentic experiences of agency are best construed as thetic states, then it seems we can have our cake and eat it. If experiences of agency are indeed thetic states, what would be special about first-person agentic knowledge and account for its asymmetry with third-person knowledge may not so much be the intentional contents of the agentic experiences on which this agentic knowledge is based as their two-faced intentional structure. What would make agentic experience yield knowledge and not just purported knowledge is the fact that predictive and retrospective inferential processes work together to ensure that our agentic experiences track the actual relations between what we try to achieve and what happens in the outside world. This is not to say that agentic experiences are an infallible source of knowledge of our own agency. Agentic experiences may sometimes be illusory. However, when they are not illusory, they can yield the kind of agentic knowledge that supporters of the middle course are looking for.

## References

- Anscombe, E. (1957). *Intention*. Oxford: Blackwell.
- Aarts, H., Custers, R., & Marien, H. (2009). Priming and authorship ascription: When nonconscious goals turn into conscious experiences of self-agency. *Journal of Personality and Social Psychology*, 96, 967-979.
- Aarts, H., Custers, R., & Wegner, D. (2005). On the inference of personal authorship: Enhancing experienced agency by priming effect information. *Consciousness and Cognition*, 14(3), 439-458.
- Bayne, T. (2008). The phenomenology of agency. *Philosophy Compass*, 3(1), 182-202.
- Bayne, T. (2009). The Sense of Agency. In F. Macpherson (Ed.), *The Senses*. Oxford: Oxford University Press.
- Bayne, T. (2010). Agentive Experiences as Pushmi-Pullyu representations. In J. Aguilar, A. Buckareff and K. Frankish (eds.) *New Waves in the Philosophy of Action*. Palgrave Macmillan, pp. 219-36.
- Bayne, T., & Pacherie, E. (2007). Narrators and comparators: The architecture of agentive self-awareness. *Synthese*, 159(3), 475-491.
- Belayachi, S., & van der Linden, M. (2009). Level of agency in sub-clinical checking. *Consciousness and Cognition*, 18, 293-299.
- Blakemore, S. J., Wolpert, D., & Frith, C. (2000). Why can't you tickle yourself? *Neuroreport*, 11(11), R11-R16.
- Blakemore, S., & Frith, C. (2003). Self-awareness and action. *Current Opinion in Neurobiology*, 13(2), 219-224.
- Blakemore, S., Wolpert, D., & Frith, C. (1998). Central cancellation of self-produced tickle sensation. *Nature neuroscience*, 1, 635-640.
- Carruthers, G. (In Press). The case for the comparator model as an explanation of the sense of agency and its break-downs. *Consciousness and Cognition*, doi:10.1016/j.concog.2010.08.005
- Desmurget, M. & Sirigu, A. (2009). A parietal-premotor network for movement intention and motor awareness. *Consciousness and Cognition*, 13, 10: 411-419.
- Desmurget, M., Reilly, K. T., Richard, N., Szathmari, A., Mottolese, C. & Sirigu, A. (2009). Movement Intention After Parietal Cortex Stimulation in Humans. *Science*, 324: 811-813.
- Dijksterhuis, A., Preston, J., Wegner, D. M., & Aarts, H. (2008). Effects of subliminal priming of self and God on self-attribution of authorship for events. *Journal of Experimental Social Psychology*, 44, 2-9.
- Eilan, N., & Roessler, J. (2003). Agency and self-awareness: mechanisms and epistemology. In J. Roessler & N. Eilan (Eds.), *Agency and self-awareness* (pp. 1-47). Oxford: Oxford University Press.

- Engbert, K., & Wohlschläger, A. (2007). Intentions and expectations in temporal binding. *Consciousness and Cognition*, 16 (2): 255-264.
- Fletcher, P. C., & Frith, C. D. (2009). Perceiving is believing: a Bayesian approach to explaining the positive symptoms of schizophrenia. *Nature Neuroscience*, 16: 48-58.
- Fourneret, P., & Jeannerod, M. (1998). Limited conscious monitoring of motor performance in normal subjects. *Neuropsychologia*, 36(11), 1133-1140.
- Frith, C., Blakemore, S., & Wolpert, D. (2000a). Abnormalities in the awareness and control of action. *Philos Trans R Soc Lond B Biol Sci*, 355(1404), 1771-1788.
- Frith, C., Blakemore, S., & Wolpert, D. (2000b). Explaining the symptoms of schizophrenia: abnormalities in the awareness of action. *Brain Research Reviews*, 31(2-3), 357-363.
- Gallagher, S. (2007). The natural philosophy of agency. *Philosophy Compass*, 2(2), 347-357.
- Haggard, P., & Clark, S. (2003). Intentional action: conscious experience and neural prediction. *Consciousness and Cognition*, 12(4), 695-707.
- Haggard, P & Eimer, M. (1999). On the relation between brain potentials and the awareness of voluntary movements. *Experimental Brain Research*, 126: 128-133.
- Haggard, P. & Magno, E. (1999). Localizing awareness of action with Transcranial Magnetic Stimulation. *Experimental Brain Research*, 127: 102-107.
- Haggard, P., Clark, S., & Kalogeras, J. (2002). Voluntary action and conscious awareness. *Nature Neuroscience*, 5(4), 382-385.
- Hendricks, K. V., Wiggers, P., Jonker, C. M., & Haselager, W. F. (2007). Towards a computation model of the self-attribution of agency. In P. Oliver, C. Kray (Eds.), *Proceedings of the artificial intelligence and simulation of behaviour annual convention* (pp. 350-356).
- Humberstone, I. L. (1992). Direction of fit. *Mind*, 101, 401: 59-83.
- Jeannerod, M. (1995). Mental imagery in the motor context. *Neuropsychologia*, 33, 11: 1419-1432.
- Jeannerod, M. (1997). *The cognitive neuroscience of action*. Oxford: Blackwell Publishers.
- Jeannerod, M. (2009). The sense of agency and its disturbances in schizophrenia: a reappraisal. *Experimental Brain Research*, 192(3), 527-532.
- Knoblich, G., & Kircher, T. T. J. (2004). Deceiving oneself about being in control: Conscious detection of changes in visuomotor coupling. *Journal of Experimental Psychology-Human Perception and Performance*, 30(4), 657-666.
- Knoblich, G., & Repp, B. H. (2009). Inferring agency from sound. *Cognition*, 111(2), 248-262.
- Knoblich, G., Stottmeister, F., & Kircher, T. (2004). Self-monitoring in patients with schizophrenia. *Psychological Medicine*, 34(08), 1561-1569.
- Korsgaard, C. 1996. *Creating the Kingdom of Ends*. Cambridge: Cambridge University Press.
- Lau, H. C., Rogers, R. D., & Passingham, R. E. (2007). Manipulating the experienced onset of intention after action execution. *Journal of Cognitive Neuroscience*, 19, 1-10.
- Leube, D., Knoblich, G., Erb, M., Grodd, W., Bartels, M., & Kircher, T. (2003). The neural correlates of perceiving one's own movements. *Neuroimage*, 20(4), 2084-2090.

- Libet, B. (1985). Unconscious cerebral initiative and the role of conscious will in voluntary action. *Behavioral and Brain Sciences*, 8: 529-566.
- Libet, B., Gleason, C. A., Wright, E. W. & Pearl, D. K. (1983). Time of conscious intention to act in relation to onset of cerebral activities (readiness potential): the unconscious initiation of a freely voluntary act. *Brain*, 106: 623-642.
- McDowell, J. (2010). What is the content of an intention in action? *Ratio*, 23: 415–432.
- McDowell, J. (2011). Some remarks on intentions in actions. *The Amherst Lecture in Philosophy* 6: 1-18. <<http://www.amherstlecture.org/mcdowell2011/>>
- Metcalf, J., & Greene, M. (2007). Metacognition of agency. *Journal of Experimental Psychology: General*, 136(2), 184–199.
- Millikan, R. G. (1995). Pushmi-pullyu representations. *Philosophical Perspectives*, IX, 185-200.
- Moore, J. W., & Fletcher, P. C. (In Press). Sense of agency in health and disease: A review of cue integration approaches. *Consciousness and Cognition* (2011), doi:10.1016/j.concog.2011.08.010.
- Moore, J. W., & Haggard, P. (2008). Awareness of action: Inference and prediction. *Consciousness and Cognition*, 17, 136–144.
- Moore, J. W., Lagnado, D., Deal, D. C., & Haggard, P. (2009). Feelings of control: Contingency determines experience of action. *Cognition*, 110(2), 279–283.
- Moore, J. W., & Obhi, S. S. (In Press). Intentional Binding and the sense of agency: A review. *Consciousness and Cognition* (2012), doi:10.1016/j.concog.2011.12.002.
- Moore, J. W., Wegner, D. M., Haggard, P. (2009). Modulating the sense of agency with external cues. *Consciousness and Cognition*, 18, 1056-64.
- Moran, R. (2004). Anscombe on "Practical Knowledge". In J. Hyman & H. Steward (Eds.) *Agency and Action*, Cambridge: Cambridge University Press, pp. 43-68.
- Naccache, L., Dehaene, S., Cohen, L., Habert, M.-O., Guichart-Gomez, E., Galanaud, D. & Willer, J.-C. (2005). Effortless control: executive attention and conscious feeling of mental effort are dissociable. *Neuropsychologia*, 43: 1318-1328.
- Pacherie, E. (2008). The phenomenology of action: A conceptual framework. *Cognition*, 107(1), 179-217.
- Pacherie, E. (2010). Self-Agency. In S. Gallagher (Ed.), *The Oxford Handbook of the Self*. Oxford: Oxford University Press, pp. 440-462.
- Pacherie, E. (2011). Nonconceptual representations for action and the limits of intentional control. *Social Psychology*, 42(1): 67-73.
- Proust, J. (2003). Perceiving intentions. In J. Roessler & N. Eilan (Eds.), *Agency and self-awareness* (pp. 296-320). Oxford: Oxford University Press.
- Sato, A. (2009). Both motor prediction and conceptual congruency between preview and action-effect contribute to explicit judgment of agency. *Cognition*, 110(1), 74-83.
- Sato, A., & Yasuda, A. (2005). Illusion of sense of self-agency: discrepancy between the predicted and actual sensory consequences of actions modulates the sense of self-agency, but not the sense of self-ownership. *Cognition*, 94(3), 241-255.

- Searle, J. (1983). *Intentionality, an essay in the philosophy of mind*. Cambridge: Cambridge University Press.
- Sirigu, A., Daprati, E., Ciancia, S., Giraux, P., Nighoghossian, N., Posada, A. & Haggard, P. (2004). Altered awareness of voluntary action after damage to the parietal cortex. *Nature Neuroscience*, 7, 1: 80-84.
- 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.
- van den Bos, E., & Jeannerod, M. (2002). Sense of body and sense of action both contribute to self-recognition. *Cognition*, 85(2), 177-187.
- Vallacher, R. R., & Wegner, D. M. (1985). *A theory of action identification*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Vallacher, R. R., & Wegner, D. M. (1987). What do people think they're doing? Action identification and human behavior. *Psychological Review*, 94, 3-15.
- Vallacher, R. R., & Wegner, D. M. (1989). Levels of personal agency: Individual variation in action identification. *Journal of Personality and Social Psychology*, 57, 660-671.
- van der Weiden, A, Aarts, H. & Ruys, K. I. (2010). Reflecting on the action or its outcome: Behavior representation level modulates outcome-priming effects on self-agency experiences. *Consciousness and Cognition*, 19, 21-32.
- Vosgerau, G., & Synofzik, M. (In Press) Weighting models and weighting factors. *Consciousness and Cognition*, doi:10.1016/j.concog.2011.09.016
- Wegner, D. (2002). *The illusion of conscious will*: The MIT Press.
- Wegner, D. M., & Wheatley, T. (1999). Apparent mental causation - Sources of the experience of will. *American Psychologist*, 54(7), 480-492.
- Wegner, D. M., Sparrow, B., & Winerman, L. (2004). Vicarious agency: Experiencing control over the movements of others. *Journal of Personality and Social Psychology*, 86(6), 838-848.
- Wong, H. Y. (In Press). A measure of my agency? *Consciousness and Cognition*, doi:10.1016/j.concog.2011.08.017