

Multiple aspects in the sense of agency¹

Shaun Gallagher

Philosophy and Cognitive Science
University of Central Florida and
University of Hertfordshire
gallaghr@mail.ucf.edu

Abstract: Recent significant research in a number of disciplines centers around the concept of the sense of agency. Because many of these studies cut across disciplinary lines there is good reason to seek a clear consensus on what ‘sense of agency’ means. In this paper I indicate some complexities that this consensus might have to deal with. I also highlight an important phenomenological distinction that needs to be considered in any discussion of the sense of agency, regardless of how it gets defined, and I offer some critical comments on some current models tied to the notion that the sense of agency has an ambiguous phenomenology.

Keywords: sense of agency, F-intention, P-intention, intentional aspect, motor control, phenomenology.

There has been a significant amount of research on and discussion of the concept of the sense of agency in recent psychology (e.g., Aartsa, Custersa, and Wegner 2005; Marcel 2003; Sato and Yasuda 2005), in philosophy (e.g., de Vignemont and Fourneret 2004; Eilan and Roessler 2003; Gallagher 2000a, 2004; Pacherie 2007; Synofzik, Vosgerau and Newen, in press), in neuroscience (e.g., Chaminade and Decety 2002; Blakemore, Frith & Wolpert 1999; Farrer et al. 2003; Tsakiris and Haggard 2005a), and in psychopathology (e.g., Daprati et al. 1997; Franck et al. 2001; Frith, Blakemore & Wolpert 2000; Stephens and Graham 2000). Because many of these studies cut across disciplinary lines there is good reason to seek a clear consensus on what ‘sense of agency’ means. In this paper I intend to show how complicated this consensus might be. A second aim is to highlight an important phenomenological distinction that needs to be considered in any discussion of the sense of agency, regardless of how it gets defined. Finally, in the penultimate section, I offer some critical comments on contemporary models that fail to take into account the phenomenological ambiguity involved in the sense of agency.

The phenomenological distinction is one that philosophers make between pre-reflective (or non-reflective) and reflective aspects of self-consciousness. Reflective self-

¹ My thanks to the special editors of this issue and to Thor Grünbaum for helpful comments. Research for this paper was supported in part by a National Science Foundation grant (# 0639037) and the European Science Foundation project *Consciousness in a Natural and Cultural Context* (BASIC).

consciousness can be further distinguished into what may be called ‘introspective reflection’ and ‘situated reflection’. Introspective reflection involves a metacognitive stance in which the subject takes her own thoughts and actions as objects for consideration. The subject might reflect on whether she is taking the right strategy to accomplish her goal, or she might ask whether what she intends to do is consistent with her beliefs, desires, and other activities. This kind of reflection may be relatively detached from current action. Situated reflection, in contrast, is embedded in an ongoing contextualized action. It involves the type of activity that I engage in when someone asks me what I am doing or what I plan to do (Gallagher and Marcel 1999). I do not necessarily frame my answers to such questions in terms of beliefs, desires, or strategies; rather I may reference the immediate environment and what needs to get accomplished.

Pre-reflective self-awareness refers to the fact that in any experience there is an implicit, first-person awareness of that experience. When I perceive or remember or imagine X, I am also aware that I am living through or undergoing this experience (see, e.g., Gallagher and Zahavi 2005; 2008; Legrand 2006; 2007). My first-order experience includes a basic feeling that I am living through that experience. This pre-reflective self-awareness, however, is a recessive or marginal rather than an explicit aspect of consciousness. When we perceive or remember, or imagine, or think about X, the central focus is on X, but it is always accompanied by a pre-reflective sense of my own experience. Alvin Goldman (1970, 96) describes this as follows.

[Consider] the case of thinking about x or attending to x. In the process of thinking about x there is already an implicit awareness that one is thinking about x. There is no need for reflection here, for taking a step back from thinking about x in order to examine it. When we are thinking about x, the mind is focused on x, not on our thinking of x. Nevertheless, the process of thinking about x carries with it a non-reflective self-awareness.

Before seeing how the distinction between reflective and pre-reflective self-awareness also applies to our actions and to the sense of agency, we need to consider some complications introduced in connection with the concept of intention.

Intentions

Some actions are deliberated about and planned out in advance. For example, I may plan to go shopping at a certain store several hours before I have the opportunity to do so. Other actions are *ad hoc* and often elicited by circumstances. For example, there is a knock on the door and I get up to answer it. Many *ad hoc* actions may be included in the framework of a larger planned action. For example, in order to initiate my shopping activity I have to get up and leave my office, which involves walking to the door and opening it. The motor components of walking to the door and opening it may be much the same whether I am on my way to the store or answering the knock. In both cases the action of walking to the door and opening it is an intentional action. Despite the similarity and perhaps the identity of the actions at the motor level, however, these are two different actions at the level of intentions. If I stop you as you walk to the door and ask what you are doing, your report should differ if you pitch it at the most pragmatic level of description: ‘I’m going shopping’, versus ‘I’m answering the knock’. The

difference comes at the pragmatic or intentional level, and it is tied to the goal or purpose of your action. But it is also the case that different intentions are cashed out at the motor level, and for different intentions even if they involve similar movements, there may be no strict identity at the motor level. Intentions reach all the way down to the motor details of the actions to the extent that the motor details support the intentions. The motor specifics of my grasp differ if I am picking up a cup in order to drink from it versus picking up the same cup in order to throw it (Jeannerod 1997; also see Pacherie 2005). Or if I want to move to the door, there are several ways to do so – I could walk or run, for example. Which one I do will depend in part on my intention and in part on the situation. If I want to lock the door because I think the knock signals the presence of danger, I may run to it; but if I am expecting someone important, my gate may be slowed so as not to appear overly eager. Slowing my gate or speeding it up may not be things that I consciously decide, but they may very well be things that are determined by circumstances (physical and social environments) as well as by the nature of my intention, whether my intention itself is well deliberated, or closer to *ad hoc*.

Actions may be more or less complicated, not simply in terms of what I am trying to accomplish (shopping, answering a knock), but also in terms of how these different aspects (planning, initiation, motor control) interconnect. In most cases, action comes along with an experience of agency, and the question I want to explore in this paper is whether the sense of agency takes on the level of complexity that corresponds to its action.

Elizabeth Pacherie (2006; 2007), with specific reference to the phenomenology of agency, recently offered a model of how the various aspects of action interconnect. She identifies three cascading “stages” of action specification. The first corresponds to the formation of future-directed intentions (F-intentions); the second corresponds to present-directed intentions (P-intentions). This distinction follows Searle’s (1983) distinction between ‘prior intentions’ and ‘intentions-in-action’, or Bratman’s (1987) distinction between future-directed and present-directed intentions. Pacherie, however, introduces a third concept: motor intentions (M-intentions).

F-intentions are formed before the action and represent the whole action as a unit. They are usually detached from the situation of action and specify types of actions rather than tokens. Their content is therefore conceptual and descriptive. F-intentions are also ... subject to distinctive normative pressures for consistency and coherence: in particular, they should be means-end coherent, consistent with the agent's beliefs and consistent with other intentions he or she may have. P-intentions serve to implement action plans inherited from F-intentions. They anchor the action plan both in time and in the situation of action and thus effect a transformation of the descriptive contents of the action plan into perceptual-actional contents constrained by the present spatial as well as non-spatial characteristics of the agent, the target of the action, and the surrounding context. The final stage in action specification involves the transformation of the perceptual-actional contents of P-intentions into sensorimotor representations (M-intentions) through a precise specification of the spatial and temporal characteristics of the constituent elements of the selected motor program. (Pacherie 2007, 3).

These three stages of intention are considered to be contributors to the experience of agency. In specific, Pacherie claims that the sense of agency contains a variety of aspects: an experience of intentional causation, the sense of initiation and the sense of control, which is further analyzable into more specific experiences. This complexity may be surprising in light of what is usually considered to be the “thin” phenomenology associated with agency, which means that the sense of agency is short-lived and phenomenologically recessive (i.e., it remains in the pre-reflective background of experience and so not very noticeable in ordinary experience), and therefore difficult to specify (e.g., Libet 1985; see Haggard 2005).

The fact that the experience of agency is phenomenologically thin, however, doesn't mean that it is not complex. There are in fact several contributories to the sense of agency and in pathological cases it may be possible for subjects to have different experiences of agency (or lack of agency) depending on whether one or another of these contributories is disrupted. Just how complex is the sense of agency, what aspects of action or action planning contribute to it, and what actually manifests itself in the phenomenology? Pacherie's analysis provides a good starting point because we can ask whether the sense of agency necessarily includes F, P, and M components, and whether there are other aspects of action that need to be considered.

Reflective and pre-reflective experiences of agency

The sense of agency (SA) is the experience that I am the one who is causing or generating the action (Gallagher 2000a&b). SA can be distinguished from the sense of ownership (SO) for movement, which is the sense that I am the one who is undergoing the movement – that it is my body moving, whether the movement is voluntary or involuntary. In the case of involuntary movement, SA is missing, but I still have SO – the sense that I am the one moving. Described in this way, these experiences are pre-reflective, which means that they neither are equivalent to nor depend on the subject taking an introspective reflective attitude, nor that the subject engages in an explicit perceptual monitoring of bodily movements. Rather, SA, like SO, is generated at the level of primary sensory-motor processes that involve efferent and afferent processes such as motor (efferent) commands and proprioception, and is manifested at the level of first-order phenomenal consciousness (Gallagher 2000a; 2005; Haggard 2005). Like my own bodily movements, my agency is not normally something that I attend to or something of which I am explicitly aware. As such, SA is phenomenologically recessive.

The senses of agency and ownership, on this description, contribute to a basic or minimal self-awareness and thereby constitute aspects of the minimal (or core or ecological) self (see e.g., Damasio 1999; David et al. 2006; Gallagher 2000, Neisser 1988). I'll refer to SA, so conceived, as SA1, the minimal or first-order sense of agency, to distinguish it from a higher-order phenomenon described by Graham and Stephens (1994; Stephens and Graham 2000). They work out an account of introspective alienation in schizophrenic symptoms of delusions of control in terms of two kinds of self-attribution.

—*Attributions of subjectivity*: the subject reflectively realizes and is able to report that he is moving. For example, he can say, "This is my body that is moving."

—*Attributions of agency*: the subject reflectively realizes and is able to report that he is the cause or author of his movement. For example, he can say "I am causing this action."

According to Graham and Stephens the SA originates at this higher-order level of attribution. Following Dennett and Flanagan, they propose an explanation of SA in terms of "our proclivity for constructing self-referential narratives" which allow us to explain our behavior retrospectively: "such explanations amount to a sort of theory of the person's agency or intentional psychology" (1994, p. 101; Stephens and Graham, 2000, p. 161). Their distinction is worked out in the context of explaining the schizophrenic symptom of inserted thought. They suggest that if we take thinking to be an action on our part, then SA for that thinking action derives from a reflective attitude toward it.

[W]hether I take myself to be the agent of a mental episode depends upon whether I take the occurrence of this episode to be explicable in terms of my underlying intentional states (1994, p. 93).

On this view we reflectively explain our actions in terms of our beliefs and desires. Accordingly, if a subject does or thinks something for which she has no intentions, and her actions fails to accord with her beliefs and desires – mental states that would normally explain or rationalize such actions – then such movements or thoughts would not appear as something she intentionally does or thinks. Whether I count something as my action thus "depends upon whether I take myself to have beliefs and desires of the sort that would rationalize its occurrence in me. If my theory of myself ascribes to me the relevant intentional states, I unproblematically regard this episode as my action. If not, then I must either revise my picture of my intentional states or refuse to acknowledge the episode as my doing" (1994, 102). On this approach, I have SA for my actions because I have a properly ordered set of second-order interpretations.

[T]he subject's sense of agency regarding her thoughts likewise depends on her belief that these mental episodes are expressions of her intentional states. That is, whether the subject regards an episode of thinking occurring in her psychological history as something she does, as her mental action, depends on whether she finds its occurrence explicable in terms of her theory or story of her own underlying intentional states. (Graham and Stephens 1994, 102; see Stephens and Graham 2000, 162ff).

According to Graham and Stephens, then, SA results from an inference made on the basis of higher-order reflective (introspective) observations. Accordingly, I'll refer to this as a reflective higher-order sense of agency (SA2), and distinguish it from the pre-reflective, first-order (minimal) sense of agency (SA1) that is phenomenologically recessive.

Reflective aspects in the sense of agency

I want to suggest that SA2 correlates with what Pacherie calls F-intentions and in some

cases, with P-intentions.² F-intentions may result from relatively abstract and conceptual considerations, or may be formulated spontaneously, as the result of situated reflection.³ In either case, the F-intention occurs sometime prior to the action. For example, considering the age of my car, my beliefs about what I can afford, and my desires for the newest models, I might decide that I will buy a new car next spring. It is equally possible that I make this decision more spontaneously when my car breaks down. Next spring when I go to the auto dealership and start looking at cars, however, I form more specific P-intentions because I have to start getting precise about what to look for. In agreement with Graham and Stephens, Pacherie indicates that F-intentions are subject to normative pressures for consistency and coherence relative to the agent's beliefs and other intentions. For example, I may reflectively deliberate about whether I should buy a car for my daughter, and I may decide to do so for a variety of reasons that are consistent with my beliefs about my daughter's good sense of responsibility and her excellent driving skills. In making the decision I've set the stage for my later action, when I actually take my daughter to the car dealer to buy the car. The claim is that the prior decision – the formation of my F-intention – adds to the sense that I am the agent of the later action. If, in contrast, without prior deliberation (without the formation of an F-intention) I stop at a car dealer to buy a car, I would not deny my agency for this action, even if I did it spontaneously. This more *ad hoc* action which involves forming a P-intention likely involves a more situated reflection or pragmatic deliberation and is, just as much as an F-intention, subject to normative pressures. Indeed, I might feel the need to explain that on seeing the red Mustang convertible in the lot, it called out to me as perfect for my daughter, or whatever. It seems that both F-intentions and P-intentions contribute to SA2 even if they do not contribute the entire sense.⁴

It's important to note that although F-intentions do add something to SA2, they are neither a necessary nor sufficient condition for SA2, since I may do something without going through a stage of deliberation or forming a prior intention, and still feel that I am the agent of that action, as Pacherie (2007, 7) points out. Moreover, and obviously, if the only thing I did were to thoughtfully deliberate and decide to do X, that would not be sufficient for generating SA for doing X, since I might never get around to

² The distinction between F-intentions and P-intentions is a relative one that depends not only on time difference between deliberation and action, but on the level of abstractness on which one decides. Clearly an F-intention is involved if in January I decide that I will buy a car in May. Now, on the first day of May I get up in the morning and I decide that today is the day to buy the car, without any specification as to what car I will buy. That too is an F-intention. But what if I don't have that thought in the morning and later in the day, on passing a car dealership and seeing a red Mustang convertible, I decide that right now is the time to buy that specific car. This seems more like a P-intention both in terms of specificity and temporal proximity to the action. But if right now I decide that I will buy that car next week, this seems more like the F-intention, despite the specificity.

³ To be clear, one could form an F-intention without thoughtful deliberation. Relative to the P-intention, however, the formation of an F-intention (deciding to do something tomorrow, for example) could itself be considered something that stands in for deliberation.

⁴ Without losing the sense of agency for an action, it also seems possible to do without any F- or P-intentions, and indeed, even do better. Dijksterhuis et al. (2006) have shown that more unconscious (or less consciously deliberative) strategies deliver better results than more conscious deliberation when deciding to buy a car or in complex cases that call for decision. Although they did not ask about sense of control or agency in the contrasting cases, there was no indication that at the point of decision people felt less in control using unconscious strategies. It has also been shown that conscious evaluative deliberation about the same object can lead to less consistent choices over time (Levine, Halberstadt and Goldstone 1996).

doing it. P-intentions seem more central to SA2, since they are more immediately tied to action context. P-intentions may depend on more immediate and situated reflections that help me to implement action; but they may also be formed spontaneously. P-intentions are not simply thoughts about what I want to do, but involve a dynamic monitoring of the action itself, a situated and pragmatically reflective verification that my action is on track.

Pace Pacherie (2007), P-intentions are not necessarily “inherited from F-intentions” since I may act without prior planning or forming a prior intention. Pacherie ties her description of P-intentions too closely to the execution of F-intention action plans, or more generally considers these intentions to be arranged in “stages.”⁵ Indeed, I would suggest that in most cases of intentional action, I engage in such action without having an action plan, and without having formulated F-intentions. For example, if I go to the kitchen to get a drink, this is certainly more *ad hoc* than planned out. Pacherie is closer to the mark when she characterizes P-intentions as “constrained by the present spatial as well as non-spatial characteristics of the agent, the target of the action, and the surrounding context” (2007, 3)

Here’s an example. I often receive e-mail from colleagues or students who ask me to write them a letter of recommendation. If circumstances are right (nothing else is pressing on my schedule) I find it easier to write the letter immediately, since I am sitting in front of my computer, and the request is right there in front of me, and otherwise I have a bad habit of forgetting such requests. My intention to write the letter is formed, and one might say it is in formation as I engage in the action of bringing up my word processing program and I start looking for information about the addressee. This action is accompanied by the sense that I could at any point forego this activity and do it later, but once started I usually continue until I arrive at the signature line and am ready to print out the letter.

In this case there seems to be no F-intention – there is no prior intention formation. One might object, however, that the F-intention was formed when I deliberated about my practice of writing letters of recommendation and decided to write them immediately upon receiving a request. But I never did do this kind of deliberation, and it is not very clear if there was one point where I formed this intention. Rather, in one instance when I received a request I actually did respond to it immediately and was satisfied that this was a better practice than forgetting to do the letter. The next time I received a request I did the same thing, and over time I built up a habitual practice that seems to guide my behavior in most circumstances. If I am following a rule here, the rule seems to have emerged from my practice, rather than the other way around. I could now formulate that rule as an F-intention, but I would be doing so only in retrospective reflection. Moreover, it would be beside the point since my habitual practice is already in place and that, rather than some F-intention, seems to be what is guiding my action. At most the late-formed F-intention could act as confirmation or reinforcement of my practice, and in that way strengthen SA.

The immediate action of writing the letter is definitely “constrained by the present spatial as well as non-spatial characteristics of the agent, the target of the action, and the surrounding context.” The sense that I could put this off, or that I could continue with it later if I were interrupted, along with the sense that I am making progress as I continue to

⁵ Pacherie (2006), however, does note that F- and/or P-intentions are not required for action, and that P-intentions may even be motivated by M-intentions.

write, and the sense of completion as I reach the point of hitting the print command – these seem to be part of what Pacherie calls the P-intention. Furthermore, I may consciously think about what I am doing in a way that keeps me doing it or puts me back on track. If I am interrupted by a phone call, as I finish the call I may ask myself what I was doing and decide to continue. It may also be the case that on returning to my computer I see the half-written letter and simply start in on it again – no thinking about it required. Of course I am likely putting some kind of thought into the formulation of the letter, but that is both part of the action and part of what keeps the action on track. There is also a wider context that comes into this action, which is indicated by the question of whether I am qualified to write the letter. This is something that I don't think about unless I suddenly feel that I am not qualified, or perhaps do not know the person sufficiently well. If the request for a letter comes from a neurosurgeon I will certainly start to reflect on whether I am qualified to write such a letter, and this reflection will reflect on just those issues that Graham and Stephens identify as contributing to SA2. That is, in such a case, I will reflectively consider the action in terms of a network of beliefs that I have about myself. But I find myself doing this only in the exceptional case when I have an unusual request – not in the case of everyday run-of-the-mill requests that call upon my known qualifications.

The P-intention, then, is closely tied to the action itself and may involve a situated reflective monitoring of the action that keeps it on track. It is contingent on the specific circumstances of action. The reflective aspects of the P-intention, however, may never arise if the action in which I am engaged is a well-practiced action, and everything runs smoothly. Indeed, well-practiced intentional action may not involve P-intentions at all. We can make this claim more precise. Pacherie lists four functions that can be attributed to P-intentions. “They trigger or initiate the intended action (*initiating function*), they sustain it until completion (*sustaining function*), they guide its unfolding (*guiding function*) and monitor its effects (*monitoring function*)” (2006, 146). In some cases, however, once initiated, an action may be carried along to completion without much sustaining effort; in other cases, sustaining effort may be necessary. Pacherie points out, correctly I think, that guiding and monitoring are control functions, and that some of these, or some part of these functions, may be taken over by M-intentions. In the case of easily accomplished, well-practiced or habitual actions environmental factors may elicit and sustain the action (think of pushing a branch out of the way of your walking path, or shaking hands with a friend), and the control mechanisms might be fully delivered at the level of M-intentions.

Intentional aspects and motor intentionality

Some intentional actions are the result of prior deliberations and F-intentions; others happen without deliberation or the formation of F-intentions. It also seems possible to engage in intentional actions and to have a sense of agency for such actions (immersed SA [Pacherie 2007] or minimal SA [Marcel 2003]) without forming P-intentions. What I have in mind here are skillful actions that are performed without being subsidiary to reflective thought; actions that involve what Dreyfus calls “spontaneous, transparent coping” (2002, 417). Such actions do require motor control and a kind of

motor intentionality, or what Pacherie calls M-intentions.⁶ A good example is driving a car. I may form an F-intention to drive rather than to take the train to work. When I start driving, this is intentional action on my part, but without any plan or explicit self-monitoring, i.e., without P-intentions that are specified to the action of driving, or even reaching my workplace. I nonetheless have a first order sense of agency (SA1) for what I am doing – I know what I am doing (although I may not know precisely how I am doing it), and I experience myself as the agent who is doing it.

Some theorists argue that SA1 is generated in motor control processes associated with bodily movement; others suggest that it belongs more properly to the realm of the intentionality (goal, purpose) of intentional action. It likely involves both (see Gallagher 2007a; in press; Haggard 2005). To see how SA1 involves motor control, consider the nature of involuntary movement. In the case of involuntary movement there is SO for the movement but no SA, or more precisely, no sense of self-agency. Awareness of my involuntary movement comes from afferent sensory-feedback (visual and proprioceptive/kinaesthetic information that tells me that I'm moving). There are no initial motor commands (no efferent signals) that I issue to generate the movement. It seems reasonable to think that in both involuntary and voluntary movement SO is generated by sensory feedback, and that in the case of voluntary movement SA1 might be generated by efferent signals that send motor commands to the muscle system. Tsakiris and Haggard (2005; also see Tsakiris 2005) review empirical evidence to support this division of labor, but also suggest that efferent processes underlying SA1 also modulate sensory feedback resulting from movement. Sensory suppression experiments (Tsakiris and Haggard 2003) suggest that SA1 arises at an early efferent stage in the initiation of action and that awareness of the initiation of my own action depends on central signals, which precede actual bodily movement. Experiments with subjects who lack proprioception but still experience a sense of effort reinforce this conclusion (Lafargue, Paillard, Lamarre, & Sirigu 2003; see Marcel 2003). As Tsakiris and Haggard put it,

the sense of agency involves a strong *efferent* component, because actions are centrally generated. The sense of ownership involves a strong *afferent* component, because the content of body awareness originates mostly by the plurality of multisensory peripheral signals. We do not normally experience the efferent and afferent components separately. Instead, we have a general awareness of our body that involves both components. (Tsakiris and Haggard 2005, 387).

Does conceiving of SA1 as generated in efferent or motor-control related brain processes provide a full account of SA1? Haggard points to another element important for SA1: “actions seem to aim towards a goal, as if pulled teleologically from the intention through to the intended effect” (2005, 292). SA1, in addition to a sense of embodied movement, involves a sense of controlling events in the external world.

⁶ One question is whether they necessarily involve M-intentions as the latter concept is defined by Pacherie: sensorimotor representations that involve a precise specification of the spatial and temporal characteristics of the constituent elements of the selected motor program. What is questionable here is whether there are representations involved at this level. The question of representationalism is something that I will set aside for purposes of this paper (but see Gallagher 2008 for a non-representational dynamic systems approach to this issue).

Haggard thus distinguishes between “urge” and “effect” – the former associated with an experience generated in motor control processes, the latter associated with perceptual monitoring of what one actually accomplishes by the action. Both of these aspects, the *intentional aspect* (what gets accomplished, or fails to get accomplished, by the action) and the *motor (or efferent) aspect* (the sense that I am causing or controlling my bodily movement) enter into SA1 (Gallagher 2007a, in press). As I drive to work I have both an occurrent sense of bodily control over the car, and I (constantly or inconstantly) monitor the traffic and road and have the sense that I am moving along in the right direction, and so forth. These aspects, and SA1 more generally, remain pre-reflective in so far as neither of them are things that I reflectively dwell upon, and indeed, as I arrive at my office I have forgotten most, if not all, of the details involved in my driving.

The distinction between these two aspects, however, is not always kept clear in the experimental literature, as a quick review of several brain imaging experiments on SA1 will show. For example, in a PET study by Chaminade and Decety (2002), subjects moved a joystick to control an icon on a computer screen in order to accomplish one of two tasks.

Task A (Leader): Subjects move their own circle icon and observe another subject’s circle following it.

Task B (Follower): Subjects follow another subject’s icon with their own.

The authors suggest that SA1 depends on what is happening on the screen (the intentional aspect) rather than what is happening in the movement of the joystick (the motor aspect, which they associate with SO).

Investigation of the neural basis of agency was performed using a paradigm in which the subjects either led (A) or followed (B) the other, in a computerized environment free of explicit reference to body parts. ... The sense of ownership, related to motor control, and the sense of agency, related to the intentional aspect, can be segregated in the analysis. (Chaminade and Decety 2002, 1977)

Chaminade and Decety associate SA1 with the intentional aspect of action – the accomplishment of the task. They nonetheless point to brain areas closely associated with motor control and body awareness as correlated with SA1. Specifically, they identify activation in the pre-supplemental motor area and the right inferior parietal cortex as responsible for generating the sense of self-agency in the leader condition (A), and in contrast, activation of the left inferior parietal cortex and the right pre-central gyrus as responsible for attributing agency to another, which they associate with the follower condition (B). One objection, of course, is that in both cases (A and B) the subject may experience SA1 for the intentional aspect – i.e., accomplishing the task. If one associates SA1 with motor control, or with brain areas responsible for motor control, then the fact that the subject intentionally moves the joystick in each task would suggest that they experience agency for each condition.

In an fMRI experiment by Farrer and Frith (2002), designed to find the neural

correlates of SA1, subjects are asked to manipulate a joystick to drive a colored circle moving on a screen to specific locations on the screen. In some instances the subject causes this movement and in others the experimenter or computer does. The subject has to discriminate self-agency and other-agency. Like Chaminade and Decety, Farrer and Frith cite the distinction between SA and SO, but associate SA with the intentional aspect of action, i.e., whether I am having some kind of effect with respect to the goal or intentional task (or what happens on the screen). Accordingly, the claim is that SO (my hand is moving the joystick) remains constant while SA1 (based on the intentional aspect) changes. When subjects feel that they are not controlling the events on the screen, there is activation in the right inferior parietal cortex (rIPC) and supposedly no SA1 for the intentional aspect of the action. When the subject does have SA1 for the action on the screen, the anterior insula is activated bilaterally.

Although Farrer and Frith clearly think of SA1 as something tied to the intentional aspect of action and not to mere bodily movement, when it comes to *explaining why* the anterior insula should be involved in generating SA1, they frame the explanation in terms of motor control.

Why should the parietal lobe have a special role in attributing actions to others while the anterior insula is concerned with attributing actions to the self? The sense of agency (i.e., being aware of causing an action) occurs in the context of a body moving in time and space. Damasio (1999) has suggested that the sense of agency critically depends upon the experience of such a body. There is evidence that both the inferior parietal lobe and the anterior insula are representations of the body the anterior insula, in interaction with limbic structures, is also involved in the representation of body schema One aspect of the experience of agency that we feel when we move our bodies through space is the close correspondence between many different sensory signals. In particular there will be a correspondence between three kinds of signal: somatosensory signals directly consequent upon our movements, visual and auditory signals that may result indirectly from our movements, and last, the corollary discharge [efferent signal] associated with motor commands that generated the movements. A close correspondence between all these signals helps to give us a sense of agency. (601-02).

Farrer et al. (2003) is a third study that sets out to discover the neural correlates of SA1. In this experiment subjects provide a report on their experience; however, all questions about agency were focused on bodily movement rather than intentional aspect. In fact, subjects were not given an intentional task to carry out other than making random movements using a joystick, and the focus of their attention was directed towards a virtual (computer image) hand that either did or did not represent their own hand movements, although at varying degrees of rotation relative to true position of the subject's hand. That is, they moved their own hand, but saw a virtual hand projected on screen at veridical or non-veridical angles to their own hand; the virtual hand was either under their control, or not. Subjects were asked about their experience of agency for control of the virtual hand movements. The less the subject felt in control, the higher the level of activation in the rIPC, and this is consistent with Farrer and Frith (2002). The

more the subject felt in control, the higher the level of activation in the right *posterior* insula, which is in contrast with Farrer and Frith (2002) where SA1 was associated with activation of the right *anterior* insula. Referencing this difference, Farrer et al. state: “We have no explanation as to why the localization of the activated areas differ in these studies, except that we know that these two regions are densely and reciprocally connected” (2003, p. 331). One possible explanation, however, is that the shift of focus from accomplishing a computer screen task (in Farrer and Frith) to controlling bodily movement (in Farrer et al.) might change the phenomenon that is being studied, even if in both cases SA1 is the supposed phenomenon.⁷

It would be helpful for all of these experiments to clearly distinguish between the intentional aspect and the motor (efferent) aspect of agency, and to say that there are at least these two contributories to SA1. The intentional aspect of SA1 is not necessarily tied to a monitoring of action at the level of P-intention, if we understand P-intentions, as Pacherie does, to involve conceptual thought, rational control, or conscious judgments.⁸ Rather, keeping track of what I am doing as I am doing some task, in many cases involves nothing more than a pre-reflective (and often perhaps a non-conscious) perceptual monitoring of what my action is accomplishing in the world. In simple actions, for example, as I reach to grasp a cup in order to take a drink, there is no conscious judgment required – that is, I don’t have to constantly verify that my reach is going in the right direction, or that my hand is shaping the proper grasp for the cup.⁹ But this doesn’t mean that I don’t perceptually monitor the location, shape, size, weight, etc. of the cup, and the success (or failure) of my grasp, at some level, or that the actual accomplishment of getting a drink, or failing to get the drink, doesn’t enter into SA1. Visual, tactile, and in this case even gustatory sensory information enters into the mix and provides information about what I am doing or what I have accomplished (or failed to accomplish) by my action. If, for instance, I bring the cup to my mouth and take a sip, but there turns out to be nothing in the cup, I will still have a sense of agency for my

⁷ There is also some suggestion that activation in IPC may be associated with SO rather than SA. Farrer et al. state: ‘Lesions of the inferior parietal cortex, especially on the right side, have been associated with delusions about the patient’s limb that may be perceived as an alien object or as belonging to another person’ (2003, 329). Such delusions are about ownership rather than agency. In addition, in Tsakiris and Haggard (2005), activity in the insula was found in the absence of movement, which implies that this area may in fact reflect body-ownership rather than agency. Farrer et al., however, suggest the following explanation. ‘We have proposed the activity seen in inferior parietal cortex relates to the feeling of loss of agency associated with the discrepancy between intended actions and sensory feedback. However, from the experiment discussed so far it is possible that the activity in this region relates solely to the sensory discordance. The feeling of agency might relate to activity in other regions. We think this is unlikely on the basis of various pathological cases in which the primary disorder concerns the feeling of agency rather than sensory discordance’ (Farrer et al. 2003, 329). But if the pathological studies they cite (i.e., alien limb) indicate something about delusions of ownership rather than agency, then this may not be so unlikely.

⁸ “The agent exercises rational control over her action insofar as (1) she is in a position to judge whether or not his way of accomplishing her action is likely to lead to success and adjusts it so as to maximize her chances of success (tracking control) and (2) she is also in a position to judge whether or not it brings about undesirable side-effects and corrects it accordingly (collateral control)... P-intentions are responsible for high-level forms of guidance and monitoring - they are concerned with aspects of the situation of action and of the activity of the agent that are both consciously perceived and conceptualized. Therefore the time scale of P-intentions is the time scale of conscious perception and rational thought” (Pacherie 2006, 150).

⁹ Exceptions to this can be found in cases of deafferentation, as studies of IW and GL have shown (see Cole 1996; Gallagher and Cole 1996; Cole and Paillard 1995).

action, but I will also be aware of a certain failure in my action, and this would contrast to the feeling of agency that I have when I satisfy my thirst.

For SA1, however, this intentional aspect amounts to nothing if it is not closely connected with the motor aspect of action which adds efferent information to the mix. On some models of schizophrenic delusions of control, for example, mechanisms associated with efference copy are disrupted and the subject accordingly lacks SA1 for his own action (Frith 1992; see Gallagher 2004). For the subject, there is still the proprioceptive sense that his hand is reaching for the cup, and all the sensory information associated with the intentional aspect is generated, but, absent efference copy which tells the motor system that it is generating the movement, the action seems alien to the subject. The subject's report is that it is his hand (that he is the one) that (who) is reaching and grasping the cup (there is SO for the movement) but he has no feeling that he has issued the motor command and so has no SA1 for the action. If that motor aspect is missing, the intentional aspect is not phenomenologically anchored in the embodied agent.¹⁰

Ambiguities in the sense of agency

Pacherie suggests that F- and P-intentions are typically conscious processes. P-intentions involve “supporting control processes responsible for keeping track of the way the agent accomplishes her action and adjusts it so as to maximize her chances of success (tracking control) and to minimize undesirable side effects (collateral control)” (2007, 4). In contrast, M-intentions often happen below the threshold of consciousness. If we include both intentional aspects and motor-control aspects under the heading of M-intentions, then it is certainly the case that intentional actions involve a large amount of non-conscious processing. This non-conscious processing, however, cashed out in neuronal terms, or in terms of motor-control models, is what generates SA1, which is a conscious experience, albeit a pre-reflective and recessive aspect of consciousness. I have also suggested that we engage in some intentional action without having SA2, without the formation of F- and/or P-intentions. Writing an academic paper at my computer, I may suddenly feel thirsty and reach for the bottle of water on my desk. I do this without formulating any kind of F- or P-intention. One might say that I do it in a close to non-conscious manner (not, of course, in the sense that I lapse into an unconscious state as I do it, but in the sense that my conscious attention continues to be focused on other things, i.e., thinking about the subject matter of the paper that I am writing). It would be better to say that I am pre-reflectively aware that I am reaching for the bottle and that this is part of what it means to have a sense of agency (SA1) for my action. If you stop my action and ask what I am doing, I could easily tell you that I am getting a drink. I will not be able to tell you, however, on the basis of my current action, how I am doing it, or what intentional aspects I am (non-consciously) monitoring, or what precise muscles I am using to accomplish the movement. There is no conscious P-intention that requires a conscious tracking of the situation. But all of the non-conscious processing that goes into

¹⁰ Citing good evidence, Synofzik, Vosgerau and Newen (in press), maintain that “When subjects plan, monitor and perceive their own actions and the corresponding effects, they often do not primarily represent them in motor-related terms (e.g. their spatio-temporal pattern), but in intentional and perceptual terms (e.g. their underlying goals).” Nonetheless it would be a mistake to completely leave aside the motor control mechanisms when explaining the SA and focus only on the intentional and perceptual goal representation.

the M-intention nonetheless generates the pre-reflective awareness that we are calling SA1.

The nature of the non-conscious processing of the M-intention can be partially specified in neuronal terms. A more complete functional model can also be specified in terms of motor-control mechanisms that involve forward and inverse models as explicated in the work of various theorists such as Wolpert (1997; Wolpert & Ghahramani 2000; Wolpert, Ghahramani & Jordan 1995), Jeannerod (1997) and Frith (1992; Frith, Blakemore & Wolpert 2000), as cited by Pacherie. Both Frith (1992) and Pacherie (2007) suggest that the forward and inverse models of motor-control can be used to explain the control of higher-order cognition. Frith uses this strategy to explain how schizophrenic thought insertion can be explained in the same terms as delusions of control that involve motor action, since both symptoms involve loss of the sense of agency. Pacherie suggests that mechanisms analogous to motor control mechanisms can explain the formation of F- and P-intentions.

The contents represented at the level of F-intentions as well as the format in which these contents are represented and the computational processes that operate on them are obviously rather different from the contents, representational formats and computational processes operating at the level of M-intentions. Yet, the general idea that internal models divide into inverse models which compute the means towards a given goal and forward models which compute the consequences of implementing these means retains its validity at the level of F-intentions. ... Similarly, it is highly plausible that action-specification at the level of P-intentions makes use of internal models ... (2007, 4).

I have elsewhere offered critical comments on Frith's proposal to model control of thinking on motor-control mechanisms (Gallagher 2004), but it is not clear that the same critique applies to Pacherie's proposal, which remains on a very general level. That our deliberation about future actions involves thinking about the means and ends of our actions seems uncontroversial. Pacherie's proposal does raise one question, however. If we regard thinking itself as a kind of action, then do we also have a sense of agency for the thinking or deliberation involved in the formation of F-intentions. It seems reasonable to suggest that if I engage in a reflectively conscious process of deliberating about my future actions and make some decisions on this basis, I would have a sense of agency for this deliberation. You could interrupt me during this process and ask what I am doing, and I could say "I'm sitting here deliberating about buying a car." The sense of agency that I feel for my deliberation process, however, is SA1. Although it is possible that I could experience SA2 for my deliberation process (I may have formed the F-intention yesterday to do my deliberations (i.e., to form my F-intentions) about car buying today), it is clearly the case that not all forming of F-intentions require SA2, otherwise we would have an infinite regress. We would have to deliberate about deliberating about deliberating, etc. Furthermore, it would seem odd to suggest that we have P-intentions for the action of forming F-intentions, unless we think of P-intentions in this case as a form of metacognition where we may be conscious of our cognitive strategies as we form our F-intentions. Certainly, however, it is not always the case that

we engage in this kind of metacognition as we formulate our F-intentions. Again, it seems that we can have a first-order sense of agency (SA1) without having SA2, if SA2 requires either prior deliberation or occurrent metacognitive monitoring.

This motivates a further question. Is SA2 for action X connected with SA1 for the action of forming an F-intention to do X? On the one hand they are obviously not equivalent since there are two different actions here: action X (e.g., buying the car), and the earlier action of deliberating to do X. On the other hand, it seems likely that the sense of agency (SA1) for my deliberation may contribute to my reflective sense that I am the agent of my own actions (SA2 for those actions).¹¹ So SA1 for forming my F-intention may contribute to SA2 for the action that follows from my F-intention. Furthermore, if I lacked SA1 for my deliberation process, that deliberation process might feel more like an intuition or an unbidden thought, or indeed, if I were schizophrenic, it might feel like an inserted thought. In any case, it might feel less than integrated with what Graham and Stephens call the “theory or story of [the subject’s] own underlying intentional states.” This, however, is precisely what they indicate to be the basis for the higher-order reflective sense of agency (SA2) for having that thought. So again it seems that SA1 for the deliberation process itself, contributes to the higher-order, reflective SA2, for forming the F-intention. Still, as I indicated in the previous section, there need not be (and, under threat of infinite regress, there can not be) a deliberation process for every action that I engage in. If in some cases there is a deliberation process, SA1 for that deliberation contributes to both SA2 for the deliberation and SA2 for the action that follows from that deliberation.

Similar questions can be raised about P-intentions. If action monitoring, at the level of P-intentions, is itself a kind of action (if, for example, it involves the act of judging), is there SA1 for that action monitoring? The processes that make up a P-intention are much closer to the intended action itself and may not feel like an additional or separate action. But if, as I have suggested, a P-intention is different from the intentional aspect of the M-intention insofar as the former is a conscious monitoring of the action, then it seems possible that I could experience SA1 for just that monitoring. I can imagine a very explicit kind of P-intention in the form of a conscious monitoring of what I am doing, for example, if I am putting together a piece of furniture by following a set of instructions. In that case I could have a sense of agency (SA1) for following the instructions and closely monitoring my actions in terms of means-ends (and certainly doing it that way would feel very different from doing it without following the set of instructions). But the SA1 for following the instructions would really go hand in glove with SA1 for the action of assembling the furniture. How we distinguish such things would really depend on how we define the action. In the process of assembling the furniture, I may start by reading instruction #1; I then turn to the pieces of wood in front of me and join two of them together. I can distinguish the act of reading from the act of joining and define SA1 for each of them. In that case, however, one can ask whether SA1

¹¹ Or, in Pacherie’s terms, the long-term sense of agency – “a sense of oneself as an agent apart from any particular action, i.e. a sense of one’s capacity for action over time, and a form of self-narrative where one’s past actions and projected future actions are given a general coherence and unified through a set of overarching goals, motivations, projects and general lines of conduct” (2007, 6) – may enter into the occurrent sense of agency for a particular action, and specifically as part of the SA2 for that action.

for the act of reading isn't just what we mean by SA2 for the act of joining. I might, however, think of the reading and the joining as one larger action of assembling the furniture, and SA1 might be defined broadly to incorporate all aspects of that assembling. It might also be the case that when I put together the second piece of furniture, I don't consult the instructions at all, in which case SA1 is more concentrated in the joining. In most practiced actions, however, a P-intention is really unnecessary because the intentional aspect of the M-intention can do the job, i.e., can keep my action on track. I might simply make up my mind (an F-intention) to do this task, and I go and immediately start to do the task without further monitoring in terms of means-ends. All of this suggests that how we discriminate SA2 and SA1 is relative to the way we define specific actions, and how practiced those actions are.

The point I want to make here is that there is some serious ambiguity, not simply in the way we define the sense of agency, but in the sense of agency itself. That is, the phenomenological ambiguity – the very ambiguity of our experience of agency – should be included in our considerations about the sense of agency. Clear-cut and unambiguous definitions make for a neat conceptual map; but the landscape itself may not be so neat. It is not always the case that P-intentions serve to implement action plans inherited from F-intentions, since there are not always F-intentions. It is not always the case that “the final stage in action specification involves the transformation of the perceptual-actional contents of P-intentions into sensorimotor representations (M-intentions) through a precise specification of the spatial and temporal characteristics of the constituent elements of the selected motor program” (Pacherie 2007, 3), since there are not always P-intentions.

Pacherie suggests that a sense of action initiation and a sense of control are “crucial” components in SA (2007, 17-18). Moreover, in both components P-intention plays a large role. But the fact that some actions for which we have SA1 take place without P-intentions puts this idea in question.

The *sense of action initiation*, Pacherie suggests, is based on the binding of P-intention and awareness of movement onset in the very small timeframe of 80-200 ms prior to actual movement onset corresponding to the time of the lateralized readiness potential, a signal that corresponds to selection of a specific motor program (Libet 1985; Haggard 2003). Specifically, she associates the P-intention with what Haggard distinguishes as urge to move and reference forward to the goal of the action. But these aspects of action experience can be, as I suggested above, pre-reflective, generated by motor-control processes, and form part of the M-intention. In this regard, it is important to distinguish P-intention from the pre-reflective perceptual monitoring of the intentional aspects of the action that can occur without a formed P-intention, as in practiced action. Whereas monitoring of the intentional aspects contribute to SA1 whether we have a conscious intention in terms of specific goals or not (Aartsa, Custersa, and Wegner 2005), the P-intention does not seem crucial for SA1.

The *sense of control* has, according to Pacherie, three dimensions corresponding to F-intentions, P-intentions, and M-intentions. Again, however, the sense of control may be reflectively conscious for F- and P-intentions, but, as generated in motor-control mechanisms, remains pre-reflectively conscious as long as the action is going well, e.g., as long as I don't stumble over or knock into something. A conscious judgment or conscious sense of control associated with the P-intention may in fact be absent until that

point when something starts to go wrong at the motor-control level.

These various conceptual distinctions – “awareness of a goal, awareness of an intention to act, awareness of initiation of action, awareness of movements, 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” (Pacherie 2007, 6) – may not show up in the actual first-order phenomenology. They may be the product of reflection on the first-order phenomenology, or the product of theoretical considerations. As I engage in action, for example, I may not experience a difference between the goal of the action and the means I use to attain the goal, although I can certainly make that distinction in my reflective consideration of my action. That distinction may show up clearly at the level of my F-intention, but may be entirely lost in my immersed SA1. My awareness of what I am doing and that I am doing it is usually struck at the most pragmatic level of description (“I’m getting a drink”) rather than at a level that distinguishes between the action and my agency, or within the action between the goal and the means, or within agency between intentional causation, initiation, and control – distinctions that Pacherie suggests can be found in the phenomenology. Phenomenologically, however, there is no such thing as a “naked intention” – the awareness of an action without an awareness of who the agent is (Jeannerod and Pacherie 2004) – or “agent-neutral” action experience (Pacherie 2007, 16). The awareness that I am the agent of an action is implicit in the pre-reflective awareness of acting, which does not contain an awareness of causation separate from awareness of control. Pacherie is thus absolutely right to note that a conceptual analysis cannot “preempt the question whether these various aspects are dissociable or not, for instance whether we can be aware of what we are doing independently of an awareness of how we’re doing it or whether we can be aware of what we are doing without at the same time experiencing this action as ours” (2007, 7). What can decide the issue, however, is agreement on where to draw the lines between *phenomenological* analysis (i.e., of what we actually experience), *neuroscientific* analysis (which may find a much finer grain of articulations at the neuronal level than show up in phenomenology), and *conceptual* analysis (which may introduce distinctions that are in neither the phenomenology nor the neurology, but may have a productive role to play in constructing cognitive models or, in regard to the individual, explaining psychological motivations, etc.).

For example, it makes good conceptual sense to distinguish between the action and the agent of the action. On some neuroscientific models, it also makes good sense to distinguish between subpersonal processes that are responsible for the perception of action and subpersonal processes responsible for the attribution of agency (Georgieff & Jeannerod 1998). But on the phenomenological level, that is, in the experience itself, when I see John open the door, I do not primarily see the action of opening the door and then secondarily attribute that action to John. I see John-opening-the-door. Likewise, when I open the door, I do not experience opening the door, and then discover that it is I who opened the door. Rather, my experience is one of I-am-opening-the-door, and clear-cut distinctions between agency and action come up only on the occasion of reflecting on the event.¹²

¹² Jeannerod and Pacherie (2004) claim that “it is like this with the perception of intention: when Mary watches John open the door, she is primarily aware of an intention to open the door, rather than being

Pathologies, of course, can throw some light on these issues. Discussions of the sense of agency are often motivated by considerations about the loss of a sense of self-agency in schizophrenic symptoms or other disorders. Delusions of control, for example, help to show the difference between SO and SA. When the patient complains that his hand is reaching for something but that he is not the one causing this movement, he is acknowledging that it is his hand and that it is moving, and thus has a sense of ownership for that movement (it's him moving rather than anyone else), but lacks a sense of self-agency for it. The subject attributes the agency of his action to someone else. But even here it is not that the subject experiences an action and then has to figure out who the agent is. Rather, the subject experiences an alien action – an action that, with regard to the experience of agency, feels to belong to someone else (Gallagher 2007a&b; in press).¹³

Conclusion

The sense of agency is both complex and ambiguous. It has multiple contributories, some of which are reflectively conscious, some of which are pre-reflectively conscious, and some of which are non-conscious. Consistent with phenomenological theories of embodiment, in everyday engaged action afferent or sensory-feedback signals are attenuated, implying a recessive consciousness of the body in action (see e.g., Gallagher 2005; Tsakiris and Haggard 2005). We do not attend to our bodily movements in most actions. We do not stare at our own hands as we decide to use them; we do not look at our feet as we walk, we do not attend to our arm movements as we engage the joystick. Most of motor control and body schematic processes are non-conscious and automatic. Just such processes nonetheless contribute to a conscious sense of agency by generating a pre-reflective awareness of our actions. Without this basic feeling of the embodied nature of action our sense of agency would be very different. In addition, we also experience, pre-reflectively, a form of *intentional feedback*, which is not afferent feedback about our bodily movements, but a perceptual sense that my action is having an effect in the world (or at least on the computer screen). This effect is not something that we reflectively dwell on, or even retain in memory. A good example of this is our usual perceptual awareness while driving a car.

The sense of agency for some actions may amount to nothing more than this (SA1). For other actions, however, the sense of agency is not reducible to just these embodied and pre-reflective processes. In addition, in many cases we may be reflectively conscious of and concerned about what we are doing. For such actions our sense of agency (SA2) will be tied to a more reflective sense of intention, by which our attention is directed toward the project or task that we are engaged in, or toward the means and/or end that we aim for.

I suggest, then, that the sense of agency, at the first-order level of experience (SA1), is already complex because it is the product of several contributory elements: efferent

primarily aware that John intends to open the door. Similarly, when Mary herself intends to open the door, she is primarily aware of an intention to open the door, rather than being primarily aware that she herself intends to open the door. Let us call this awareness of an unattributed or 'naked' intention" (Jeannerod and Pacherie 2004, 116).

¹³ It feels that it does belong to someone else, not simply "as if" it belonged to someone else. The latter is the experience in Anarchic Hand Syndrome.

signals, sensory (afferent) feedback, and intentional (perceptual) feedback. If any of these contributory elements fail, or fail to be properly integrated, then we can get a disruption or disturbance in the sense of agency. SA1 can be supplemented and complicated by SA2, which is based on higher-order considerations about whether what I plan to do or have done is consistent with my belief system, or with my conception of efficient means-end relations.

Although conceptually we may distinguish between different levels (first-order, higher-order), and neuroscientifically we may be able to identify different brain processes responsible for these different contributories, in our everyday phenomenology we tend to experience agency in a more holistic, qualitative, and ambiguous experience that might be open to a description in terms of degrees of control.

More conceptually and generally, we can identify four different contributories to the sense of agency.

- Basic efferent motor-control processes generate one aspect of SA1 as a first-order experience linked to bodily movement in and towards an environment.
- Pre-reflective perceptual monitoring of the effect of my action in the world, which generates a second aspect of SA1, the first-order experience of the intentional aspect of action.
- Formation of F-intentions, that often involve the reflective deliberation or planning that precedes action and which contributes to SA2, a second-order, reflective attribution of agency in terms of a sense of consistency with personal beliefs, etc.
- Formation of P-intentions, that is, the conscious monitoring of action in terms of specific means-ends relations, further contributing to a more specific SA2 in terms of a higher-order control.

In any particular case the sense of agency might be disrupted in different or multiple ways depending on what contributory element is disrupted. The attribution of self-agency (SA2) may be disrupted by problems with introspective higher-order cognition (Stephens and Graham 2000) or formation of F-intentions (Pacherie 2007). A good example of this is the case of narcotic addiction, as discussed by Frankfurt (1988). If a drug addict invests himself in resisting drugs he may feel that something other than himself is compelling him to drug use. If he withdraws from taking the drug, when he starts using again he may not conceive of himself as the agent.

It is in virtue of this identification and withdrawal, accomplished through the formation of second-order volition, that the unwilling addict may meaningfully make the analytically puzzling statements that the force moving him to take the drug is a force other than his own, and that it is not of his own free will but rather against his will that this force moves him to take it (Frankfurt 1988, 18; see Grünbaum, in press, for discussion).

At a different level the sense of agency may be disrupted by problems with motor control mechanisms – a failure of efference signals (Tsakiris and Haggard 2005) or the integration of sensory and motor signals (Farrer et al. 2003). It may also be disrupted by

a lack of concordance between P-intention and the intentional aspects associated with the effects of action. Thus, the loss of the sense of agency in various pathologies – including schizophrenia, anarchic hand syndrome, obsessive-compulsive behavior, narcotic addiction, etc. – may in fact be different sorts of loss.

References

- Aartsa, H. Custersa, R. Wegner, D. M. 2005. On the inference of personal authorship: Enhancing experienced agency by priming effect information. *Consciousness and Cognition* 14: 439–458
- Blakemore, S., Frith, C., & Wolpert, D. (1999). Spatiotemporal prediction modulates the perception of self produced stimuli. *Journal of Cognitive Neuroscience* 11: 551–559.
- Bratman, M. E. 1987. *Intention, Plans, and Practical Reason*. Cambridge, MA: Cambridge University Press.
- Chaminade, T. and Decety, J. 2002. Leader or follower? Involvement of the inferior parietal lobule in agency. *Neuroreport* 13 (1528): 1975-78.
- Cole, J. 1995. *Pride and a Daily Marathon*. Cambridge, Massachusetts: MIT Press.
- Cole, J. and Paillard, J. 1995. Living without touch and peripheral information about body position and movement: Studies upon deafferented subjects. In J. Bermudez, A. Marcel, and N. Eilan (eds.), *The Body and the Self* (pp. 245-66). Cambridge, MA: MIT Press.
- Damasio, A. 1999. *The Feeling of What Happens: Body and Emotion in the Making of Consciousness*. New York: Harcourt Brace and Co.
- Daprati, E., Franck, N., Georgieff, N., Proust, J., Pacherie, E., Dalery, J. & Jeannerod, M. (1997). Looking for the agent: an investigation into consciousness of action and self-consciousness in schizophrenic patients. *Cognition* 65 (1): 71-86.
- David, N., Bewernick, B.H., Cohen, M.X., Newen, A., Lux, S., Fink, G.R., Shah, N. J., and Vogeley, K. 2006. Neural representations of self versus other: Visual-spatial perspective taking and agency in a virtual ball-tossing game. *Journal of Cognitive Neuroscience* 18 (6): 898–910.
- Dijksterhuis, A. Bos, M. W., Nordgren, L. F., van Baaren, R. B. 2006. On making the right choice: The deliberation-without-attention effect. *Science* 311: 1005–7.
- Dreyfus, H. 2002. Intelligence without representation: Merleau-Ponty’s critique of mental representation. *Phenomenology and the Cognitive Sciences* 1 (4): 367-383.
- de Vignemont, F. and Fournieret, P. 2004. The sense of agency: A philosophical and empirical review of the “Who” system. *Consciousness and Cognition* 13: 1–19.
- Eilan, N. and Roessler, J. 2003. Agency and Self-Awareness: Mechanisms and Epistemology. In J. Roessler & N. Eilan (Eds.), *Agency and Self-Awareness* (1-47). Oxford: Oxford University Press.
- 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. *NeuroImage* 18: 324–333

- Farrer, C. and Frith, C.D. 2001. Experiencing oneself vs. another person as being the cause of an action: the neural correlates of the experience of agency. *NeuroImage* 15: 596-603.
- Franck, N., Farrer, C., Georgieff, N., Marie-Cardine, M., Dalery, J., d Amato, T., & Jeannerod, M. (2001). Defective recognition of one's own actions in patients with schizophrenia. *American Journal of Psychiatry* 158: 454–459.
- Frith, C. D. 1992. *The Cognitive Neuropsychology of Schizophrenia*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Frith, C. D., Blakemore, S., & Wolpert, D. 2000. Abnormalities in the awareness and control of action. *Philosophical Transactions of the Royal Society of London* 355: 1771–1788.
- Gallagher, S. 2000a. Philosophical conceptions of the self: implications for cognitive science. *Trends in Cognitive Science* 4 (1): 14-21.
- Gallagher, S. 2000b. Self-reference and schizophrenia: A cognitive model of immunity to error through misidentification. In D. Zahavi (ed.), *Exploring the Self: Philosophical and Psychopathological Perspectives on Self-experience* (pp. 203-39). Amsterdam & Philadelphia: John Benjamins.
- Gallagher, S. 2004. Neurocognitive models of schizophrenia: A neurophenomenological critique. *Psychopathology* 37: 8-19.
- Gallagher, S. 2005. *How the Body Shapes the Mind*. Oxford: Oxford University Press.
- Gallagher, S. 2007a. The natural philosophy of agency. *Philosophy Compass*. 2 (j.1747-9991.2006.00067.x). (<http://www.blackwell-synergy.com/doi/full/10.1111/j.1747-9991.2007.00067.x>)
- Gallagher, S. 2007b. Sense of agency and higher-order cognition: Levels of explanation for schizophrenia. *Cognitive Semiotics* 0: 32-48.
- Gallagher, S. (2008 – in press). Are minimal representations still representations? *International Journal of Philosophical Studies* 16 (3).
- Gallagher, S. (in press). Agency, free will, and psychopathology. In J. Parnas and Kenneth S. Kendler (eds.), *Philosophical issues in psychiatry: Natural kinds, mental taxonomy and the nature of reality*. Baltimore: Johns Hopkins University Press.
- Gallagher, S. and J. Cole. 1995. Body schema and body image in a deafferented subject. *Journal of Mind and Behavior* 16: 369-390.
- Gallagher, S. and Marcel, A. 1999. The self in contextualized action. *Journal of Consciousness Studies* 6 (4): 4-30.
- Gallagher, S. and Zahavi, D. 2005. In E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (<http://plato.stanford.edu/archives/sum2007/entries/self-consciousness-phenomenological/>).
- Gallagher, S. and Zahavi, D. 2008. *The Phenomenological Mind*. London: Routledge.
- Georgieff, N. and Jeannerod, M. 1998. Beyond consciousness of external events: A 'Who' system for consciousness of action and self-consciousness. *Consciousness and Cognition* 7: 465-77.
- Goldman, A. 1970. *A Theory of Human Action*. New York: Prentice-Hall.
- Graham, G. and Stephens, G. L. 1994. Mind and mine. In G. Graham and G. L. Stephens (eds.), *Philosophical Psychopathology* (pp. 91-109). Cambridge, MA: MIT Press.

- Grünbaum, T. (in press). Action and agency. In S. Gallagher and D. Schmicking (eds.), *Handbook of Phenomenology and Cognitive Science*. Dordrecht: Springer.
- Haggard, P. 2003. Conscious awareness of intention and of action. In J. Roessler & N. Eilan (Eds.) *Agency and Self-awareness*. Oxford: Oxford University Press, pp. 111-127.
- Haggard, P. 2005. Conscious intention and motor cognition. *Trends in Cognitive Sciences* 9 (6): 290-95.
- Jeannerod, M. 1997. *The Cognitive Neuroscience of Action*. Oxford: Blackwell Publishers.
- Jeannerod, M. and Pacherie, E. 2004. Agency, simulation, and self-identification. *Mind and Language* 19 (2): 113-46.
- Lafargue, G., Paillard, J. Lamarre, Y., and Sirigu, Y. 2003. Production and perception of grip force without proprioception: is there a sense of effort in deafferented subjects? *European Journal of Neuroscience* 17 (12): 2741-9.
- Legrand, D. 2006. The bodily self. The sensori-motor roots of pre-reflexive self-consciousness. *Phenomenology and the Cognitive Sciences* 5: 89-118.
- Legrand, D. (ed.) 2007. Pre-reflective self-consciousness and action. *Consciousness and Cognition*. Special Issue, 16 (3).
- Levine, G. M., Halberstadt J. B. and Goldstone, R. L. 1996. Reasoning and the weighting of attributes in attitude judgments. *Journal of Personality and Social Psychology* 70: 230-240 (1996).
- Libet, B. 1985. Unconscious cerebral initiative and the role of conscious will in voluntary action. *Behavioral and Brain Sciences*, 8: 529-66.
- Marcel, A. 2003. The sense of agency: Awareness and ownership of action. In J. Roessler & N. Eilan (Eds.), *Agency and Awareness* (48-93). Oxford: Oxford University Press.
- Neisser, U. 1988. Five kinds of self-knowledge. *Philosophical Psychology*, 1: 35-59.
- Pacherie, E. 2007. The sense of control and the sense of agency. *Psyche* 13 (1). (<http://psyche.cs.monash.edu.au/>).
- Pacherie, E. 2006. Towards a dynamic theory of intentions. In S. Pockett, W.P. Banks & S. Gallagher (eds.), *Does Consciousness Cause Behavior? An Investigation of the Nature of Volition* (145-167). Cambridge, MA: MIT Press.
- Pacherie, E. 2005. Perceiving intentions. In *A Explicação da Interpretação Humana*, João Sááguas (ed.), Lisbon: Edições Colibri, 2005, pp. 401-414.
- Sato, A. and 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: 241-255.
- Searle, J. 1983. *Intentionality: An Essay in the Philosophy of Mind*. Cambridge: Cambridge University Press.
- Stephens G. L. and Graham G. 2000. *When Self-Consciousness Breaks: Alien Voices and Inserted Thoughts*. Cambridge MA: MIT Press.
- Synofzik, M. Vosgerau, G., and Newen, A. (in press). Beyond the comparator model: A multifactorial two-step account of agency. *Consciousness and Cognition*, In Press, Corrected Proof, Available online 4 May 2007
- Tsakiris, M. 2005. On agency and body-ownership. Paper presented at Expérience Subjective Pré-Réflexive & Action (ESPRA) Conference, CREA, Paris. December 2005.

- Tsakiris, M. and Haggard, P. 2003. Awareness of somatic events associated with a voluntary action. *Experimental Brain Research* 149:439–446
- Tsakiris, M. and Haggard, P. 2005. Experimenting with the acting self. *Cognitive Neuropsychology* 22 (3/4): 387–407.
- Wolpert, D. M. 1997. Computational approaches to motor control. *Trends in Cognitive Sciences*, 1, 6: 209-216.
- Wolpert, D. M. & Ghahramani, Z. 2000. Computational principles of movement neuroscience. *Nature Neuroscience Supplement*, 3: 1212-1217.
- Wolpert, D. M., Ghahramani, Z. & Jordan, M. I. 1995. An internal model for sensorimotor integration. *Science* 269: 1880-1882.