

**teorema**

Vol. XXXVII/3, 2018, pp. 21-38

ISSN 0210-1602

[BIBLID 0210-1602 (2018) 37:3; pp. 21-38

## First Person Awareness of Agency

José Luis Bermúdez

### RESUMEN

Muchos filósofos y científicos cognitivos mantienen que los juicios de agencia se basan en el llamado *sentido de agencia*, que se considera como algo dado fenomenológicamente, parte del tejido de la experiencia. Este artículo rechaza la idea de que el sentido de agencia sea un componente irreductible de la fenomenología a nivel personal. En cambio, defiende que el sentido de agencia puede ser entendido en términos de aspectos más fundamentales de la fenomenología de la acción. El artículo evalúa diferentes propuestas para una explicación deflacionista de la conciencia de agencia a la luz de importantes hallazgos en el estudio experimental de la conciencia de la acción.

PALABRAS CLAVE: *agencia; intención, sentido de la agencia; Anscombe, integración intencional, control.*

### ABSTRACT

Many philosophers and cognitive scientists hold that judgments of agency are grounded in the so-called *sense of agency*, which is construed as a phenomenological given, part of the fabric of experience. This paper rejects the view that the sense of agency is an irreducible component of personal-level phenomenology. Instead, it argues that the sense of agency can be understood in terms of more fundamental aspects of the phenomenology of action. The paper evaluates different proposals for a deflationary account of the awareness of agency in the light of significant findings in the experimental study of action awareness.

KEYWORDS: *Agency; Intention; Sense of Agency; Anscombe; Intentional Binding, Control.*

What is first person awareness of agency? This question has been glossed in different ways in different areas of philosophy and cognitive science. One approach, originating in the work of Elizabeth Anscombe, focuses on intentional action: What sort of knowledge does an agent have of her own intentions as she acts intentionally? Another philosophical approach has explored the issue of first person awareness of agency from the more global perspective of self-consciousness: What is it to be aware of oneself as an agent? Recent discussions of agency within cogni-

tive science have focused on the distinct but related issue of how agents can discriminate actions that they themselves are carrying out from bodily movements that are not actions, on the one hand, and different agents' actions, on the other: How can I be aware, in a given situation, that it is *me* who is *acting*?

In the following I focus in particular on the first and third approaches, although the second is never far away. To focus the issue, I want to begin with a very commonsense idea. If  $\phi$  is a type of action, then to be aware of oneself as a  $\phi$ -ing agent is to be able to ascribe the act of  $\phi$ -ing to oneself. In other words, to be aware of oneself as a  $\phi$ -ing agent is to be able to judge: "I am  $\phi$ -ing". That seems uncontroversial, but of course it does not take us very far. We need to know what the basis is for the self-ascription. What are the grounds for the judgment? In virtue of what is the self-ascription made?

As far as the *self* part of the self-ascription is concerned, many cognitive scientists are committed to a common approach. They hold that judgments of agency are grounded in what is often called a sense of agency. Here is a characterization from a recent review article:

[The sense of agency] SoA is the phenomenal experience of initiating and controlling an action. It is the feeling of authorship that we refer to when we say sentences like "*I am the one who is in control of this car*" or "*It must have been me who just pressed this button*." As such, SoA phenomenally distinguishes our own self-generated actions from those actions generated by others [Braun et al. (2018) pp. 1-2].

A range of tasks, some implicit and some explicit, are standardly taken as evidence for the existence of a sense of agency. The sense of agency is construed as a phenomenological given, part of the fabric of experience. The standard view is that it is because we experience a sense of agency when we act that we are able to make judgments of agency. The sense of agency has subpersonal origins, of course. So, the standard view might be represented like this:

subpersonal mechanisms → sense of agency → judgments of agency

There is a range of different models as to what those subpersonal mechanisms consist in and how they operate.

Broadly speaking, the models fall into two broad types [see Moore (2016) for an overview]. Models of the first type are predictive. The sense of agency arises when sensory feedback matches the predictions

generated by internal forward models of action. This type of comparator model was first proposed more generally in motor control by Frith [Frith (2005)]. A second approach, certainly more popular and arguably with stronger empirical support, is retrospective. On this approach, the sense of agency emerges from a (subpersonal) process of retrospective inference, driven by consistency between an observed action and a prior intention, where the intention is the most likely explanation of the action [Wegner and Wheatley (1999)]. A multifactorial model combining contextually weighted predictive and retrospective elements has also been proposed [Synofzik et al. (2008)].

My question for this paper is not whether there is a phenomenologically salient sense of agency, but rather whether it is an irreducible component of personal-level phenomenology. That is to say, can the so-called sense of agency be understood in terms of more basic elements of the phenomenology of action? This topic has not been explicitly explored within the cognitive science literature, but it does seem to be implicitly answered there in the negative, since the theoretical models just referred to all move immediately from the personal-level sense of agency to modeling the subpersonal mechanisms that might be generating it. They leave no space for interpreting the sense of agency as grounded in other aspects of personal-level phenomenology.

I am pretty confident that Elizabeth Anscombe would not have been very sympathetic to the idea that we have a sense of agency that underwrites self-ascriptions of intentional action.<sup>1</sup> Moreover, she is insistent that judgments of agency are not grounded in aspects of the phenomenology of agency. That, I take it, follows from her view that our knowledge of our own intentional actions is, in her well-known phrase, *knowledge without observation*. The details of her positive account of knowledge without observation are not easy to decipher, but the negative dimensions of the view are clear enough.<sup>2</sup>

In *Intention* she explicitly compares an agent's knowledge of her own intentional actions with her knowledge of how her limbs are arranged. She writes:

A man usually knows the position of his limbs without observation. It is without observation, because nothing *shows* him the position of his limbs; it is not as though he were going by a tingle in his knee, which is the sign that it is bent and not straight [Anscombe (1957) §8].

Taking the analogy at face value, Anscombe is committed to denying that judgments of agency are grounded in the phenomenology of agency.

The same reasons that lead her to deny that there are phenomenological signs telling the agent how her limbs are distributed must lead her to deny that there are phenomenological signs yielding knowledge of how she is intentionally acting. In other words, judgments of agency cannot be rationally responsive to the phenomenology of agency, because knowledge without observation is not grounded in phenomenology.

I propose to explore an alternative to both the views just discussed – the standard cognitive science view, on the one hand, and Anscombe’s, on the other. On this alternative view, *contra* Anscombe, judgments of agency are grounded in, and based upon, features of our experience as acting subjects. When we make judgments of agency we do so for reasons, and those reasons are experiential reasons. When our judgments of agency express knowledge, the knowledge that they express is observational knowledge in Anscombe’s sense – that is to say, there are aspects of our personal-level experience that “shew” us that we are acting and how we are acting. (Of course, there is no suggestion that this knowledge is observational in a more standard sense of the term. It is first personal – knowledge from the inside, not knowledge derived from taking a third person perspective on oneself.)

On the other hand, *contra* the standard cognitive science view, those features of our experience on which judgments of agency are grounded have to be appropriately independent of the judgments that they ground. From an epistemological perspective, where our primary interest is in the reasons for which judgments are made and how they are justified, discussions of the sense of agency can easily seem uninformative. If we ask what the rational grounds are for our judgments of agency, then it seems unilluminating to be told that we judge that we are acting because we feel that we are acting, or that we have a sense that we are acting.

The problem is not just that talk of feelings and a sense of agency is somewhat amorphous, although that may certainly be a concern. More fundamental are concerns about independence. It is a basic requirement that any kind of evidence or ground provide independent support for what it is supposed to ground or provide evidence for.<sup>3</sup> If evidence and hypothesis are too closely linked, then the former cannot support the latter – or the latter go beyond the former.<sup>4</sup> In this case, the problem is that, as things stand, we do not have a characterization of what is to have a sense of agency that is independent of the judgments for which it is being proposed as a support. One index of this is that the sense of agency is often operationalized in terms of judgments of agency.

This emerges particularly clearly when we look at explicit experimental measures for the sense of agency. These are typically based on self-reports. So, for example, experimental subjects are shown an action on a screen at the same time as they are performing a similar action. They might be tapping their fingers and show a video of finger-tapping, which could be their own fingers or the experimenter's fingers. Gloves or some similar device are typically used to create uncertainty, and the subjects are asked whether they are seeing their own fingers move or whether it is the experimenter's fingers (as, for example, in Farrer et al. (2008), discussed further below). In other words, judgments of agency are used as evidence for the existence of a sense of agency without any clear indication of how the sense of agency and the judgment of agency might differ, or how the former could support the latter.

This is not a criticism of experimental discussions of the sense of agency, of course. From an experimental perspective, issues of evidence and rational grounds are not a primary concern. The sense of agency is operationally defined. Experimenters use "sense of agency" as a label for an effect that can be reliably elicited in various contexts, and that appears to have a fairly clearly defined functional role. The effect can be measured explicitly, as just discussed, or implicitly (using, for example, the intentional binding paradigm, discussed further below).

However, these reflections do suggest that we need to go beyond operationally defined constructs in order to illuminate the epistemology of judgments of agency and hence, by extension, for a full understanding of first person awareness of agency. But the projects are complementary, not opposed to each other, and for that reason I will use the extensive experimental literature on the sense of agency to refine and sharpen the discussion of the epistemology of agency.

So, with that in mind, in the remainder of this paper I consider three different proposals for grounding judgments of agency in the phenomenology of intentional action.

#### *Proposal 1*

Judgments of agency are grounded in joint awareness of (i) a goal-directed bodily movement, and (ii) the goal to which the bodily movement in (i) is directed.

#### *Proposal 2*

Judgments of agency are grounded in joint awareness of (i) a goal-directed bodily movement; (ii) the goal to which the bodily move-

ment in (i) is directed; and (iii) an intention to achieve the goal in (ii) through the bodily movement in (i).

*Proposal 3*

Judgments of agency are grounded in joint awareness of (i) a goal-directed bodily movement; (ii) the goal to which the bodily movement in (i) is directed; (iii) an intention to achieve the goal in (ii) through the bodily movement in (i); (iv) being able to control the bodily movement in (i) in the service of achieving the goal in (ii).

There is considerable simplification here. Much of the experimental literature on agency is focused on relatively simple actions, such as finger tapping and moving joysticks, for obvious reasons of experimental tractability. And so, the three proposals have been formulated in a way that should make it easier to map onto experimental discussions. How to scale those proposals up to more complicated actions (of the type that feature more often in philosophical discussions of action) is a matter for another occasion.

*Proposal 1*

Some philosophical accounts of action employ the concept of a basic action. Roughly speaking, the idea is that actions are hierarchically and instrumentally structured. I might simultaneously be performing a range of different actions, but those actions are not independent of each other. Agents perform some actions by performing others. On a view that goes back to Arthur Danto, this hierarchical and instrumental process bottoms out in basic actions [Danto (1963)]. A basic action is an action such that every other action that an agent might be performing at a given moment is ultimately done through performing that action. A natural suggestion is that basic actions are bodily movements. So, for example, I perform the action of setting up a meeting by performing the action of sending a group email, which I ultimately achieve through performing the basic action of tapping the keyboard in a complex sequence. A variant on the basic action view, going back to Davidson and Anscombe, holds that it is descriptions of actions, not actions themselves, that are hierarchically and instrumentally organized.<sup>5</sup> On this variant, then, what is basic is the description of my action in terms of a complex sequence of keyboard presses.

The details of the basic action view are not important for present purposes.<sup>6</sup> The main point I want to extract is the suggestion that, in the last analysis, awareness of agency will have to bottom out in awareness

of some sort of bodily movement. One reason for adopting this view is that awareness is factive. One cannot be aware that  $p$  unless  $p$  is actually the case. Applying this to awareness of agency, it follows that one cannot be aware that one is  $\phi$ -ing unless one actually is  $\phi$ -ing. But, for every non-basic action (or non-basic action description, on the variant view) a potential gap opens up between thinking that one is  $\phi$ -ing and actually  $\phi$ -ing. If the internet connection is down, then the complicated sequence of keyboard presses cannot succeed in setting up the meeting, and so I cannot be aware of setting up the meeting. But still, I can be aware of the basic bodily movements (the keyboard presses) by which I am trying to set up the meeting.<sup>7</sup>

At the same time, of course, awareness of basic bodily movements cannot be sufficient for awareness of agency. There are plenty of bodily movements that are not actions. At a minimum, actions are goal-directed. They have a purpose, which is not just a key part of what distinguishes actions from “mere” bodily movements, but also what individuates them *qua* action. So, it seems reasonable to require, first, that only goal-directed bodily movements can count as (basic) actions and, second, that awareness of agency incorporate some sort of awareness of a goal.

These two thoughts yield the background to the first proposal:

*Proposal 1*

Judgments of agency are grounded in joint awareness of (i) a goal-directed bodily movement, and (ii) the goal to which the bodily movement in (i) is directed.

At first pass, however, Proposal 1 is unlikely to capture the phenomenon fully. One thing that seems missing is awareness of the connection between the goal in (ii) and the goal-directed bodily movement in (i). It is not enough for awareness of agency that there be awareness of a goal-directed bodily movement and awareness of the goal to which that bodily movement is directed. There also needs to be awareness that the bodily movement is directed towards the goal. (I take it that we can be aware of a goal-directed bodily movement without being aware that it is directed towards a particular goal, just as we can be aware of a car traveling towards Houston without being aware that it is traveling towards Houston.)

The experiments mentioned earlier by Farrer et al. vividly illustrate the problem. The experiments were designed to investigate the neural correlates of the sense of agency. While in a scanner, subjects performed a lengthy tapping task with their middle and index fingers while wearing a glove – the task was to maintain a constant alternation between the two

fingers, following initial training with a metronome. Video of their movements over a 150 second interval was projected onto a screen with a consistent delay (of 800 or 1,000 ms). Although the video in fact showed only their own movements, subjects were told that the footage randomly and imperceptibly switched over to showing someone else performing the same task. Subjects were asked to press a button when they thought they were watching their own movements, and when they were watching someone else's movements. The experimental paradigm produced a bi-stable impression of agency, with subjects switching several times between (correctly) self-attributing the movements and (incorrectly) attributing them to someone else.<sup>8</sup>

Conditions (i) and (ii) in Proposal 1 both seem to be satisfied in the Farrer et al. experiments, both in the cases where the subjects make correct judgments of agency, and when they make incorrect ones. The subjects are aware of a goal-directed bodily movement, namely the finger tapping movements, and they were also aware of the goal (namely, to keep the two fingers alternating in a regular and constant manner). But their judgments of agency switch arbitrarily from correct to incorrect (and vice versa). Clearly, the information that they have is insufficient for first person awareness of agency, as manifested in being able to make judgments of agency that are, broadly speaking, reliable.<sup>9</sup>

As suggested earlier, what seems to be missing is awareness of the connection between the goal-directed bodily movement and the goal. The experimental subjects are aware of the goal and can evaluate whether or not the goal is being attained (because they can observe on the screen the regularity and constancy of the finger-tapping). They are also aware of their own bodily movements (via proprioception, kinesthetic awareness, and so forth). But, I hypothesize, the delay between movement and visual feedback made it hard for them to see how (and whether) their own bodily movements contributed to achieving the goal, which is why they alternate between attributing the seen movement to themselves and denying that it is theirs.

### *Proposal 2*

This brings us to Proposal 2. The second proposal closes the gap between goal-directed bodily movement and goal by incorporating awareness of the intention governing the movement:

#### *Proposal 2*

Judgments of agency are grounded in joint awareness of (i) a goal-directed bodily movement; (ii) the goal to which the bodily move-

ment in (i) is directed; and (iii) an intention to achieve the goal in (ii) through the bodily movement in (i).

Going back briefly to the Farrer et al. experiments, one of the confounding effects of introducing delayed feedback is to make it difficult for subjects to monitor the execution of their intention. So, Proposal 2 can be seen as a way of clarifying the breakdown of awareness of agency in those experiments.

Proposal 2 may seem to run into a circularity objection. After all, many philosophers have thought that intentions count as mental actions, and so to be aware of an intention is *ipso facto* to be aware of one's own agency.<sup>10</sup> How then can we include awareness of intention in an informative account of the grounds for judgments of agency?

However, awareness of "ordinary", bodily actions is a very different type of achievement from awareness of mental actions. Awareness of agency (in the non-mental sense) permits agents to attribute actions to themselves. That is why we are discussing it as a ground for judgments of agency. As we have been discussing, self-attributions of (non-mental) action can manifest the agent's ability to distinguish her *actions* from her "mere" bodily movements, and also to distinguish *her* actions from the actions of another. For this reason, the psychological role of awareness of agency is bound up with being able to ask two questions of a given piece of observed or experienced behavior. First, it makes sense to ask: is this behavior an action? And second, it makes sense to ask: is it *me* who is performing this action?

It is doubtful, I think, that the first question can reasonably be posed with respect to mental actions. I do not see how the distinction between genuine action and mere behavior can be developed with respect to mental actions. But even if it can, the second question is quite plainly inappropriate for mental actions. As many have pointed out, one of the defining features of our own introspective awareness of our own mental states (including mental actions, if such there be) is the impossibility of being aware of a mental state and wondering whether it is one's own. This is the much-discussed phenomenon of *immunity to error through misidentification relative to the first person pronoun* first broached in Shoemaker (1968). Mental actions are immune to error through misidentification in this sense, but non-mental actions are not. Much of the experimental literature on agency is based, in effect, upon presenting subjects with a version of the question: Someone is performing this action, but is it me?

For this reason, then, even if intentions are mental actions, to be aware of an intention is very different in type and kind from being aware of a non-mental act. In fact, I would suggest that awareness of a mental act does not really involve awareness of agency at all, at least not in anything like the sense in which awareness of a physical action might involve awareness of agency. In any event, there is no real sense in which Proposal 2 can plausibly be accused of appealing to the very phenomenon that it is trying to explain.

So, with that out of the way, how convincing is Proposal 2? As mentioned earlier, incorporating awareness of the intention closes the gap between the goal and the goal-directed bodily movement. The concept of an intention has of course been developed in many different ways, but it is almost always closely linked to the concepts of reason and purpose. To perform an action intentionally is to perform it for a reason and with a specific purpose in view. This teleological dimension can be more or less explicit in ways that Searle's well-known distinction between prior intentions and intentions-in-action is intended to capture [Searle (1983)]. At a minimum, therefore, to be aware of an intention is to be aware that a particular bodily movement is being performed, or will be performed, in order to achieve a specific goal.

Experiments on intentional binding provide an interesting framework for exploring Proposal 2. As mentioned earlier, the experimental literature on the sense of agency uses both implicit and explicit measures of the sense of agency. Explicit measures are self-reports, such as those in the Farrer et al. experiments. Intentional binding is probably the leading implicit measure of the sense of agency. The intentional binding phenomenon was first reported in Haggard, Clark, and Kalogeras (2002) and occurs when subjects are asked to report the perceived times of actions and their subsequent effects. In the classic paradigm, subjects are seated in front of a clock with rotating hands on a computer screen (similar to that used in the well-known Libet experiments). The baseline is set in two separate conditions, where subjects are asked to report the position of the clock hand when they either make a voluntary key press or when they hear the sound of a tone. In experimental conditions, subjects make a voluntary key press, which is invariably followed 250ms later by the sound of a tone. In some blocks, subjects were asked to report the position of the clock hand when they made the key press, while in others they were asked to report the time of the tone. The intentional binding effect is revealed when the baseline conditions are compared to the experimental conditions. The experimenters found that the perceived time

of the voluntary action was later in the experimental conditions (when the action had an apparent effect) than in the baseline condition, while conversely the perceived time of the tone was earlier in the experimental conditions (when it was apparently brought about by the subject's own action) than in the baseline condition. In other words, the perceived interval between an action and its result is compressed relative to the baseline condition.<sup>11</sup>

The occurrence of intentional binding is widely accepted as an index of the sense of agency, although there is no generally accepted account of the role of this compression of perceived time in the experience of agency [see Moore and Obhi (2012)] for a review of the main experimental research into and theoretical discussion of, intentional binding).<sup>12</sup> For present purposes the basic idea of intentional binding as an implicit measure of the sense of agency is all we need, because it gives us an alternative to self-report for identifying when subjects are aware of their own agency. In particular, we can, I think, assume the following. If intentional binding is not present, then judgments of agency are unlikely to be justified. This in turn means that where there is no intentional binding, then it seems very likely that at least some of the grounds for judgments of agency are absent.

With that in mind, we can turn to some interesting experiments reported in Haggard and Clark (2003), originally designed to arbitrate between predictive and retrospective interpretations of intentional binding (see above for the distinction in the more general context of the sense of agency). Haggard and Clark set up a standard intentional binding paradigm, as just discussed, with a specific movement of the index finger followed after 250ms by a beep. The distinctive twist in this experiment was that in key conditions the finger movement was artificially generated via TMS (transcranial magnetic stimulation). The TMS-induced movement was physically very similar to the intentionally generated movement and was also followed by a beep after 250ms.<sup>13</sup> Subjects were asked to initiate a finger movement at a point of their choosing within the first 3.5 seconds of the trial. TMS stimulation was randomly programmed in the same time interval, with the result that on some trials the intention to move the finger was "naturally" implemented, while in others it was preempted by the TMS-induced movement.

The important result here was that intentional binding occurred only on those trials where the intention was *not* preempted by the TMS-induced movement. The relevance to Proposal 2 should be clear. All three elements in Proposal 2 are present in the cases where the TMS-

induced movement intervenes between the agent's intention to produce the beep by moving her index finger and the sound of the beep. Awareness of the intention (iii) is obviously present, as is awareness of the goal. And the design of the experiment ensures that the awareness of the bodily movement is consistent across the two different types of trial.<sup>14</sup> Plainly, therefore, something is missing from Proposal 2. That brings us to the third and final proposal for an account of awareness of agency.

### *Proposal 3*

We started with a minimal account of awareness of agency in terms of joint awareness of a goal and corresponding goal-directed bodily movement. Adding awareness of the intention in Proposal 2 brought into play awareness of the connection between the bodily movement and the goal, since it is part of the content of an intention that *this* bodily movement is directed towards achieving *that* goal. But, as we have just seen, all three of those factors can be in play without awareness of agency (as measured by intentional binding). What could be missing?

Michael Bratman's influential account of intentions points us in a promising direction, I think.<sup>15</sup> For Bratman, intentions have three defining features. First, they are irreducible to beliefs and desires, and so make up an autonomous category of psychological state. Second, they are stable – that is, they tend to persist, absent new information and/or some process of rational reconsideration. Third, they are what he terms “conduct-controlling pro-attitudes” [Bratman (1987), p. 27]. I'd like to focus on this last point, on the idea that intentions control action. This is one of the ways in which he thinks that intentions differ fundamentally from desires. One can have desires without acting upon them, but to have an intention is to be ready and committed to acting upon it when the time comes.

Bratman tends to understand the idea of control in terms of action initiation, but it is certainly broader than that. Intentions continue to be controlling once action has begun, because they govern the process of execution and implementation. This is made possible by awareness of agency. Part of what it is to be aware that one is acting intentionally is to be able to monitor that action on an ongoing basis, using that monitoring as a basis for making ongoing adjustments.

Just as actions are hierarchically organized (or hierarchically described), this type of control can be exercised at multiple levels, and in response to different types of feedback. At higher levels, instrumental sub-goals can be adjusted in pursuit of an action's over-arching goal. At

lower levels, action control manifests itself in fine-grained monitoring, modification, and fine-tuning of physical movements.

Adding awareness of control to the elements in Proposal 2 yields our last proposal:

*Proposal 3*

Judgments of agency are grounded in joint awareness of (i) a goal-directed bodily movement; (ii) the goal to which the bodily movement in (i) is directed; (iii) an intention to achieve the goal in (ii) through the bodily movement in (i); (iv) being able to control the bodily movement in (i) in the service of achieving the goal in (ii).

I would suggest that awareness of component (iv) is what is disrupted when intentional actions are preempted by the TMS-induced movements discussed above. The TMS-induced movement is indeed a goal-directed bodily movement and the agent can be aware of it as such, but she is not aware of it as something that is subject to her control. It is this lack of control, I conjecture, that explains why there is no awareness of agency in these experimental conditions.

The experiments that we have been looking at have all involved very simple bodily movements (such as finger tapings) that do not typically require much by way of online motor control, let alone more complicated types of instrumental monitoring of subgoals, but there is some experimental evidence pointing to the importance of control-awareness for awareness of agency. So, for example, there have been several studies of how the sense of agency is affected by distorting visual feedback from an action. In Farrer, Frey et al. (2008), for example, subjects were asked to use a joystick to move randomly but continuously during a 70s interval. They were prevented from seeing their hands, instead watching a virtual hand moving on a screen in front of them. Experimenters were able to distort the movement of the virtual hand with respect to angle – by 25 or 50 degrees – and subjects were told (correctly) that the virtual hand was sometimes moved by their own joystick and at other times by an experimenter's joystick. The task was to identify whether what they were seeing was their own movement; a distorted version of their own movement, or someone else moving the joystick – in other words, to make judgments of agency.

The aim of the experiment was to explore the neural correlates for the sense of agency, but the important point for present purposes is that the judgments of agency became increasingly unreliable as the angular distortion increased. This, I would suggest, is because the angular distor-

tion decreases the subjects' awareness of their control over the bodily movements that (through proprioception and kinesthesia) they know themselves to be making. Awareness of control in the case of these movements (which, although temporally extended, are still very simple) is essentially awareness of being able to modify them, but these becomes increasingly hard to envisage when there is no consistent movement feedback – when, as in angular distortion cases, visual feedback and proprioceptive feedback start to diverge.

### *Conclusion*

In Bermúdez (2011) I distinguished between inflationary and deflationary conceptions of what it is to experience one's body and one's limbs as one's own – what is often referred to as the sense of ownership. Inflationary and deflationary theorists both hold that there is a positive phenomenology of ownership, but they differ in how they understand it. On inflationary views, judgments of ownership are grounded in an irreducible "feeling of mineness", whereas deflationary views ground judgments of ownership in more fundamental aspects of bodily experience.<sup>16</sup>

Discussions of the sense of ownership and the sense of agency often take place together and it is not surprising that we can find a parallel distinction between inflationary and deflationary accounts of the sense of agency. An inflationary account would be one that grounds judgments of agency in an irreducible feeling of agency, whereas a deflationary one would seek the grounds for judgments of agency in more fundamental aspects of the experience of agency.

This paper has explored three different proposals for a deflationary account of the sense of agency. The final, and to date most satisfactory, proposal holds that judgments of agency are grounded in the agent's

- (i) awareness of a goal-directed bodily movement;
- (ii) awareness of the goal to which the bodily movement in (i) is directed;
- (iii) awareness of an intention to achieve the goal in (ii) through the bodily movement in (i);
- (iv) awareness of being able to control the bodily movement in (i) in the service of achieving the goal in (ii).

The idea that judgments of agency are grounded in (i) through (iv) certainly seems to be consistent with the current experimental data, and with

plausible philosophical account of action. The next step in developing this model must surely be to scale up from very simple actions that are really just glorified bodily movements to more complex forms of intentional action. Actions further up the hierarchy of intention are much harder to tackle *directly* with experimental methods, but this is an evolving field where the dialog between theory and experiment is bound to continue.

*Department of Philosophy*  
*Texas A & M University*  
 4223 TAMU  
 College Station TX, 77483-4223 USA  
 E-mail: [jbermudez@tamu.edu](mailto:jbermudez@tamu.edu)

## NOTES

<sup>1</sup> In Bermúdez (2011) and Bermúdez (2018) I offer a broadly Anscombean argument against the view that there is a phenomenologically irreducible sense of bodily ownership (a distinctive feeling in virtue of which we experience our bodies and our limbs as our own). A structurally similar line of argument would be effective against the idea that we have a sense of agency.

<sup>2</sup> For interpretations of Anscombe's positive views on knowledge of intentions, see Haddock (2011) and Setiya (2011).

<sup>3</sup> At a minimum, we would normally expect  $1 > p(H/E) > p(H)$ . so that (i) the hypothesis is not guaranteed to hold, given the evidence, but (ii) the presence of the evidence increases the probability of the hypothesis. But obviously this leave plenty of details to be resolved.

<sup>4</sup> Interestingly, Anscombe herself was a vocal proponent of this requirement of independence. For further discussion of the requirement itself, and how Anscombe discusses it, see Bermúdez (2002) and (2018).

<sup>5</sup> See Anscombe (1957) and, for example, Davidson (1969).

<sup>6</sup> For some of the complexities in trying to make this notion precise, see the essays in Part 1 of Dancy and Sandis (2015).

<sup>7</sup> Certainly, one might reasonably think that I can be aware of trying to act in a certain way. But trying to act is not acting, and so awareness of a trying is not awareness of agency.

<sup>8</sup> The experiments revealed the angular gyrus to be a key neural area for self-attributions of agency.

<sup>9</sup> There is certainly no requirement that awareness of agency be infallible, but one would expect the grounds on which it is based to permit the experimentally demonstrated fluctuations without any corresponding change in information.

<sup>10</sup> See Geach (1957) for an early defense of mental actions, and the essays in O'Brien and Soteriou (2009) for more recent discussion.

<sup>11</sup> It's worth pointing out that no perceived intervals were actually measured. In both cases we have, as it were, a derived perceived interval, calculated from independent judgments of action and tone.

<sup>12</sup> There are some obvious questions one might ask here. Is intentional binding phenomenologically salient? That is, is there an experience of compressed time, as distinct from a compressed experience of time? If there is, is this part of the experience of agency, or simply an experiential concomitant of agency? In the first case, what are the other components of the experience of agency? In the second case, what is the direction of causation? Does the experience of agency cause the experience of compressed time, or does the direction of causality go the other way?

<sup>13</sup> To be clear on what “very similar” means here, I quote from the paper: “We took care to position the TMS coil in each subject so as to produce involuntary movements of the index finger, minimising contraction of more proximal muscles and muscles activating other digits. In both conditions, therefore, the movements selectively involved just one effector. On the other hand, the muscle activity in the two cases is necessarily different: voluntary actions produced a sustained burst of EMG activity lasting approximately 100 ms, while MEPs produced a single rapid twitch. However, this physical difference is unlikely to account for the different binding patterns of actions and MEPs, for several reasons. First, our binding estimates are based on subtracting a baseline judgement for each event occurring in isolation. This compensates for differences in perceived time due to the different physical properties of the events. Second, in a recent study, we showed that the perceived time of MEPs can be modulated along the same lines as the perceived time of other events [Tsakiris & Haggard, (2003)]. In that study, the MEP was the second event in a pair, and occurred either as the somatic effect of a voluntary keypress, or following a passive displacement of the subject's finger against the key. We showed that the MEP was bound forward in time to the voluntary action that caused it, but pushed away from the passive displacement. Thus, the perceived time of MEPs is not intrinsically any less modifiable than other events.” [Haggard and Clark (2003) p. 704].

<sup>14</sup> Again – see the passage quoted in n. 13.

<sup>15</sup> See Bratman (1987) and the essays in Bratman (1999).

<sup>16</sup> See Bermúdez (2017) for a deflationary analysis of the sense of ownership in terms of the experienced space of the body.

#### REFERENCES

- ANSCOMBE, G. E. M. (1957), *Intention*; Oxford, Basil Blackwell.
- BERMUDEZ, J. L. (2011), “Bodily Awareness and Self-Consciousness”; *Oxford Handbook of the Self*; S. Gallagher, Oxford, Oxford University Press, pp. 157-179.
- (2017), “Ownership and the Space of the Body”; *The Body and the Self Revisited*; A. Alsmith and F. de Vignemont, Cambridge MA, MIT Press, pp. 117-143.

- (2018), “Bodily Ownership, Bodily Awareness, and Knowledge Without Observation”; *The Bodily Self*, J. L. Bermudez, Cambridge MA, MIT Press, pp. 183-202.
- BRATMAN, M. (1987), *Intentions, Plans, and Practical Reason*; Cambridge MA, Harvard University Press.
- (1999), *Faces of Intention: Selected Essays on Intention and Agency*, Cambridge University Press.
- BRAUN, N., et al. (2018); “The Senses of Agency and Ownership: A Review”; *Frontiers in Psychology* 9 (535) DOI: 10.3389/fpsyg.2018.00535 [online].
- DANCY, J. and C. SANDIS (2015), *Philosophy of Action: An Anthology*; Wiley-Blackwell.
- DANTO, A. (1963); “What We Can Do”; *Journal of Philosophy* 60(15), pp. 435-445.
- DAVIDSON, D. (1969), “The Individuation of Events”; *Essays in Honor of Carl G. Hempel*, N. Rescher. Dordrecht, Reidel; pp. 216-234.
- FARRER, C. and C. D. FRITH (2002), “Experiencing Oneself vs Another Person As Being the Cause of an Action: The Neural Correlates of the Experience of Agency”; *Neuroimage* 15(3), pp. 596-603.
- FARRER, C., FREY, S. H. et al. (2008), “The Angular Gyrus Computes Action Awareness Representations”; *Cerebral Cortex* 18(2), pp. 254-261.
- FRITH, C. (2005), “The Self in Action: Lessons from Delusions of Control”; *Consciousness and Cognition* 14(4), pp. 752-770.
- GEACH, P. (1957), *Mental Acts*, Routledge and Kegan Paul.
- HADDOCK, A. (2011), “The Knowledge that a Man has of his Intentional Actions”; *Essays on Anscombe’s Intention*, A. Ford, J. Hornsby and F. Stoutland, Harvard University Press, pp. 143-169.
- HAGGARD, P., et al. (2002), “Voluntary Action and Conscious Awareness”; *Nature Neuroscience* 5, pp. 382-5.
- HAGGARD, P. and S. CLARK (2003), “Intentional Action: Conscious Experience and Neural Prediction”; *Consciousness and Cognition* 12(4), pp. 695-707.
- MOORE, J. W. (2016), “What Is the Sense of Agency and Why does it Matter?”; *Frontiers in Psychology* 7 (1272); DOI: 0.3389/fpsyg.2016.01272 [online].
- MOORE, J. W. and S. S. OBHI (2012), “Intentional Binding and the Sense of Agency: A Review”; *Consciousness and Cognition* 21(1), pp. 546-561.
- O’BRIEN, L. and M. SOTERIOU (2009), *Mental Actions*, Oxford University Press.
- SEARLE, J. R. (1983), *Intentionality: An Essay in the Philosophy of Mind*; Cambridge University Press.
- SETIYA, K. (2011) “Knowledge of Intention”; *Essays on Anscombe’s Intention*. A. Ford, J. Hornsby and F. Stoutland, Harvard University Press, pp. 170-197.
- SHOEMAKER, S. (1968), “Self-Reference and Self-Awareness”; *Journal of Philosophy* 65(19), pp. 555-567.
- SYNOFZIK, M., et al. (2008), “Beyond the Comparator Model: A Multifactorial Two-Step Account of Agency”; *Conscious Cogn* 17(1), pp. 219-239.
- TSAKIRIS, M. and P. HAGGARD (2003), “Awareness of Somatic Events Associated with a Voluntary Action”; *Experimental Brain Research* 149 (4) pp. 439-446.

WEGNER, D. M. and T. WHEATLEY (1999), “Apparent Mental Causation. Sources of The Experience of Will”, *American Psychologist* 54 (7), pp. 480-492.