Final version in Roessler, J. and Eilan, N. eds. 2003. *Agency and Self-Awareness,* Oxford: Clarendon Press, pp 358-382.

On Knowing One's Own Actions[1]

1. Introduction

Given the recent debates about self-knowledge and first-person authority it is surprising that there has not been more discussion about our knowledge of our actions.[2] It is surprising because our knowledge of our own actions seems, *prima facie,* to share many of the features of our knowledge of beliefs and perceptions, that have given rise to these debates.[3] At least this is what I want to suggest in this chapter. Indeed, perhaps the main motivation for this chapter is to place discussion of our knowledge of our own actions firmly alongside our knowledge of other psychological phenomena. I will outline what seem to be intuitively plausible features of our knowledge of our actions, and consider what account we might give of such knowledge that respects those features. A suggestion as to what form an account of our knowledge of our actions should take will be offered.

These tasks will constitute the explicit content of the chapter. However, it is also my hope that discussion of the knowledge we have of our own actions will serve some further aims. I hope it will serve to extend the range of possible sources for self-knowledge. By 'self-knowledge' I mean not just knowledge of that thing which is in fact the knower, but knowledge of ourselves *as* ourselves. The account sketched here promises to give us a way of knowing of our own actions that is unavailable to anyone other than the acting subject.

---

[1] This paper was mostly written during leave taken under the Research Leave Scheme of the Arts and Humanities Research Board. I am very grateful to them. I have benefited from very useful comments from participants of meetings in Fribourg, Girona, Manchester and Warwick (the latter in a seminar arranged as part of the AHRB Project in Consciousness and Self-Consciousness). Special thanks to the participants in a graduate research seminar on self-knowledge at UCL, and for comments to Alan Brown, David Levy, Paul Noordhof, Mark Sacks, Barry Smith and an anonymous referee for this volume. Particular thanks are due to the editors for their comments, interest and patience.

[2] There was a flurry of papers on the subject in the early sixties. See O'Shaughnessy (1963); Anscombe (1963); Donnellan (1963); Broadie (1967); Olsen (1969); Danto (1963). Since then there has been Velleman's important discussion in his *Practical Reflection*. It was only after writing this paper that I came across Johannes Roessler's valuable contribution to this issue. See Roessler (this volume).

[3] See Cassam (ed.) (1994), C. MacDonald, B. Smith & C. Wright (1998) and Shoemaker (1996) as key sources for the work that constitutes these debates.

An assumption that forms part of the background to this chapter is the assumption that actions are as primitive a psychological phenomenon as beliefs and perceptions.[4] And it aims to consider what account we should give of our actions so understood. In doing this, the chapter aims to take its place in a move away from a general *passivism* about the psychological, a move away from the idea that the basic case of a mental phenomenon is a mental *state* or *disposition*. A Martian reader of contemporary philosophy of mind would be quite likely to think that we do not really believe that we judge or act. Or at least she would think that we understand such things as psychologically non-basic: as structured groups of more basic psychological, and possibly non-psychological, states. While there are many who claim themselves to be non-reductivist about psychological phenomenon, the reduction or elimination of mental activity in favour of mental passivity is in fact widespread.

2. Features of our knowledge of our own actions.

Let us consider a simple case of action. I want to catch my friend's attention and I raise my arm in order to do so. Now consider my knowledge of my action of raising my arm. Intuitively, such knowledge is puzzlingly easy to come by. I seem to know directly and authoritatively that I am voluntarily raising my arm in a way that I do not know others' actions directly and authoritatively. Further, when I know that I am raising my arm, I seem to know it no later than when I have started to raise it. In particular, I do not seem to have to await perceptual information, for example, that muscles are contracting and that my arm is rising in order to know it. Also, wanting to know what I am doing seems to be all that is required in order for me, in normal cases, to know what I am doing. If I wonder what I am doing when I am raising my arm, I seem to be immediately supplied with an answer. My knowledge of what I am doing seems, normally, to be immediately available given that I am acting. Let us separate these features of our knowledge of our action under three headings and say that intuitively our knowledge of our actions appears to be:

First person authoritative: An agent seems to be authoritative over her own actions, and in a way that she is not over other's actions.

[4] I argue elsewhere that we should take bodily actions as prime and as not amenable to reductive analysis in terms of intentions and bodily movements. Although no reductive analysis of actions is available a more modest account of actions is.

Relatively a priori. Our knowledge of our actions appears to be spontaneous and to be given immediately upon acting. It does not seem to require any investigation or to be based on evidence.
Relatively transparent or self-intimating. Our actions, like other psychological phenomena seem to have a certain conditional epistemic availability to us. It does not seem to be the case that all our actions could as a matter of brute fact be beyond our ken.

Let us go through these features in a bit more detail:
Authoritive: I am going to take first-person authority with respect to a given subject area to imply that there is the possibility of an epistemic first/third-person asymmetry with respect to that area. Here the asymmetry must be taken to lie not simply in the fact that the subject can in central cases know more or better than others, with respect to that area, because they are around a lot or are more interested. Rather, the subject can know more or better because they know in a way which is in principle unavailable to others. So let us say:

> *X is first person authoritative with respect to a fact, p, about X iff in central cases of X's judgements concerning p, we can say that X is in a better position than others to know p, because X knows p in a way in principle unavailable to others.*

Given this understanding of first-person authority it is, I think, plausible to claim that we are authoritative with respect to our actions. This authority is exemplified by the fact that when a subject acts, and so moves their body, we take the subject to be authoritative relative to others, about whether the subject *acted* in so moving or not. Consider the case where I intentionally raise my arm. Given that there is an action of me raising my arm, it will almost always be the case that I am able to know that I raised it, and know in a way that others do not. Others will look to me to know whether I raised my arm voluntarily. If I tell them that it was a voluntary act on my part, and not an involuntary, unwilled, movement they will presume me to be right. Of course, it is possible that I am deluding myself. It is also possible that I fail to be sincere and have reasons for wanting the person to think that, contrary to the facts, I acted voluntarily. However, for the main, and in the normal case, it will be presumed

that I am in a better position to say when I am acting, and when I am not, than they are. We might not take a subject as authoritative about her action described as a movement of her body. However, as long as we keep clear the distinction between authority with respect to movements of the body (arm risings), and authority with respect to *actions* (arm raisings), then it is clear that there is a first/third person asymmetry with respect to the latter.[5]

Let me emphasize that on this understanding, first-person authority with respect to our own actions *does not mean* that it is not possible for us to be wrong about whether we acted, and that another cannot be better position than me to know whether I acted.

We are not infallible with regard to our actions and can clearly think that we acted while having failed – due to some motor failure, say – to do so. On an occasion where such a thing occurs another may be in a position to put me right. But note that we have here a parallel with perception, and possibly belief. We claim a first-person authority with respect to our own perceptions, and beliefs, even though we can think that we are perceiving something or having belief about something when we are not.

Further, given the possibility of self-deception and unconscious action, we might fail to know that we are acting when we are. In such cases, another subject who sees us moving may be better able than us to know that we are acting. My deep desire to catch someone's attention may result in my voluntarily dropping my handkerchief, but the desire may be a sufficiently uncomfortable one for me to entirely disavow the action, and for me to think that it was involuntary.

Nevertheless, it seems clear that given certain background conditions which we are entitled to take as met – such as that the subject's body is functioning normally or that there are no special repressive mechanisms in place – we can assume a first-person and third person asymmetry over the question of whether the subject acted.

[5] It should be noted that given bodily awareness – kinaesthesia in particular – it might be thought that we also have first person authority about whether our bodies moved. It seems to me that the nature of the authority attached to such cases is complex and is weaker than the phenomenon identified, but I will not discuss it further here.

Authority with respect to knowing that I acted vs. knowing what I did.

Let us suppose that it is agreed that we do *seem* to know *that* we have acted in a way that is in principle unavailable to others, and agreed that this seems to give us a certain kind of epistemic authority over our actions. It may nevertheless be said that this does not give me authority over *what* I have done. To know *that*, it may be said, I have to have recourse to my perceptual faculties. And if my knowledge of what I have done is grounded in perception, then any first/third-person asymmetry in my knowledge of what I have done will be due to the fact that only I have perceptual access through bodily awareness to the activities of my own body. It will not be special to action, but rather will be a feature of my knowledge of bodily movement more generally.

Well, I can clearly be ignorant, relative to others, about many of the things I have done. The tendency to describe our actions in terms of their effects – effects that may well be unknown to us – means that we very often can be said not to know what we have done. Given that actions can be described in terms of their consequences, and given that I can be ignorant of the consequences of my actions, I can be ignorant of my actions under such descriptions. It is perhaps mainly for this reason that we overlook the authority that we have over own actions.[6] So, it is important to admit that I did not seem to have any special authority about *what* I did when what I did comes under the description of unintended consequences. However, it does not follow from this that I did not have any authority about what I did relative to descriptions which are more basic. If there are descriptions of actions which are somehow basic, we *can* make our claim a claim about our authority over our basic actions.

One suggestion might be that if we take intentions and/or tryings as necessary for action, then there will always be a description of my actions in terms of what I

[6] Physical actions are not the only psychological phenomena we describe in terms of their consequences. We ascribe mental actions in terms of their consequences – we say things like ‘Your decision not to go to the party was a decision not to meet NN’ even when the subject being addressed does not know NN or that NN was to be at the party. Note that in a not dissimilar, but in a much more restricted way, we ascribe people beliefs in terms of their implications or pre-suppositions. ('Your belief that women are foolish is a belief that your own daughter is foolish', 'Your belief that water is wet is a belief that H_2O is wet', for example.) Perhaps, the way we very often re-describe the objects of perception, beyond any capacity the viewer has to recognize the objects as falling under those descriptions, comes closest to the action case. ('He was looking at a genuine Goya, priced at ten pounds and did not buy it'.)

intended or tried to do, and that description may be thought basic, relative to other descriptions. And, it may be said, I do seem to be authoritative about what I have done, when what I have done is described under the description drawn from what I intended or tried to do.

However, things are more complex than this suggests. Our intentions and our tryings can also appear to fall under competing descriptions. So we can say: 'You think you are intending/trying to ring the door bell of no. 6, but you are not, that's no. 4's bell you are intending/trying to ring'.[7] Perhaps we should think of these as *de re* ascriptions the truth of which like *de re* belief ascriptions, often transcend our capacity to accept them as true. The claim could then be that the statement of authority should be taken to be utilizing *de dicto* ascriptions.

I do not think that this is the right way to go. First, it is far from obvious that intentions or tryings, as causal precursors, *are* necessary for actions. It is possible, and in my view plausible, to hold that that there can be deliberate actions which are not preceded by any intention to act. One might think that intentions should be understood as effective ways of storing conclusions of practical reasoning for the future, which are not needed in cases where an action itself is the conclusion of an exercise of practical reason. It is also possible to deny that we try to act whenever we act, even when we succeed in acting. One might think that tryings to act are kinds of degenerate or failed actions. Settling these questions falls beyond anything that can be accomplished here, but their openness raises a concern with the suggestion mooted.

Second, it is implausible to say that I am in general authoritative with respect to my actions when those actions are described in terms of what I *de dicto* intended or tried to do. Consider a subject intending or trying to get a ball in the corner pocket in a game of snooker. Suppose that the subject acknowledges that this is the right way to describe what they are intending or trying to do. Imagine they strike the ball and the ball rolls into the corner pocket. While we do want to say that the subject is authoritative about whether they acted, it seems implausible to claim that the subject

[7] The example comes from O'Shaughnessy (1980, Vol 2, p.85.)

is authoritative, in contrast to others, about whether they are getting the ball in the corner pocket.

The trouble with descriptions in terms of what a subject is intending or trying to do is that they seem to avert to the subject's main purpose or motive in doing what they are doing. It seems, however, that while the subject will also be authoritative about what her purpose in acting is, and authoritative about the fact that she is acting for the purpose of doing one thing rather than another, she may not be authoritative about whether her purpose came off.

Consider, in contrast, a subject's relation to her basic actions. Basic actions are those actions a subject can carry out directly, without having to do anything else; they are the actions that a subject needs to do in order to do anything else. I think it is plausible to suppose that the descriptions which correspond to a subject's basic actions will be descriptions in terms of bodily movements, and so plausible to suppose that the subject is authoritative with respect to bodily movement descriptions such as 'raising my arm' or 'lifting my foot'. I am, however, not going to aim to settle here the question of exactly which actions are basic actions. The point I want to urge here is that, whichever actions we think are basic actions, we are justified in supposing that an agent will have a grasp of the possible ways in which they are able to act directly. For, if an agent had no such grasp of the possible ways that they could act directly, it is hard to see what sense we could make of the agent determining to act in one way, rather than another, in order to realize their aims. Given that the ability to carry out a basic action is presupposed by the ability to carry out any action, we then are justified in supposing that it is a precondition of any action that an agent has a grasp of the possible ways they can act, which are in this way basic. If the above is right, then I think we have some way of seeing how it could be that an agent can be first-person authoritative not only about the fact that they have acted but also about what they have done. The suggestion is that knowledge of *what* one is doing, and not just knowledge *that* one is acting, comes from my acting against a grasp of possible things I could have done as basic actions. It is my opting to carry out this basic action, rather than that, which, assuming I have the requisite concepts, will give me the description under which I am authoritative with respect to the action. Further, although I may be authoritative with respect to a basic action φ, and I may intend or try to do ϕ by doing

φ, it will not follow that I am authoritative with respect to my actions ϕ. It will not follow, because I may not be authoritative about whether by doing φ I succeeded in doing ϕ. So, while I might intend to pot the snooker ball by moving my arm forward, I can be authoritative about whether I moved my arm, without being authoritative about whether I potted the ball.[8]

In summary, it seems hard to deny that we are authoritative relative to others about the fact that we have acted when we have. But, I have suggested, we also have reason to think that there is an important asymmetry relative to others in our way of knowing what we have done. Given that prior to any action we must have a grasp of the possible basic ways that we might act, and given that we choose to act in one way rather than another, relative to that grasp, then we seem have the materials to account for a way of knowing what we have done that could not be available to anyone other than the agent, and which does not call on our perceptual faculties to provide immediate grounds. So, if this picture is right, the suggestion that a subject who acts is authoritative relative to others about their actions – both in terms of what we have done and of whether we have acted – becomes compelling. Further, if this picture is right, we have some explanation for the second feature identified above of our knowledge of our actions.

Relatively a priori: Perhaps the most notable and problematic feature of our knowledge of our own actions is that it appears to be immediately available to the subject who is engaged in acting, and who is aiming to answer the question as to whether they are acting. Such knowledge does not seem to have to await the testimony of our perceptual faculties.[9] We do not seem to need to feel our muscles clenching and our arm rising, or to see the trajectory of our hand, in order to know what we are doing. Perceptual knowledge is, without doubt, required as part of the background that makes action possible. However, it seems that in order to know what we are doing in the case of an individual action we do not need to perceive

[8] Note that we have here a close parallel with our authority over our perceptions. I know authoritatively that I am seeing, and I know what I am seeing, but my authority over the latter seems only to extend to descriptions in terms of basic observational concepts.

[9] *Pace* Anscombe I do not think we need to think of knowledge of our bodies through bodily awareness as non-perceptual. I therefore mean to include bodily awareness among those sources not directly necessary for knowledge of what we are doing.

simultaneously, either via bodily awareness or our other five senses, what we are doing.[10] In this way, our knowledge of our action is relatively a priori: that is, it is not acquired, or justified, via perceptual evidence in a given case, although it does rest on the obtaining of background conditions which we are entitled to take as met, and which themselves may garner support from perceptual information.

Relatively transparent: It is not just that I do not seem to have to perceive that I am acting in order to know that I am acting, nor just that I am authoritative when I do judge that I am acting. Also the subject's actions, like other psychological phenomena, seem to have a certain conditional epistemic availability to the judging subject. We can obviously fail to know that we are acting, as when we are acting absentmindedly or are repressing what we are doing. But, it does not seem to be the case that our actions can be, as a matter of brute fact, beyond our ken. It is, I think, very hard for us to imagine an agent who is capable of asking themselves the question 'What am I doing?' not being able normally to answer the question correctly. This would be to imagine agent capable of reflexive thought, voluntarily carrying out one action, rather than another, and yet not knowing that they are acting. It seems to me, however, that we cannot, in Shoemaker's phrase, envisage a creature who is simply self-blind with respect to all their actions in this way. There seems rather to be a necessary and conceptual connection between a subject acting and its knowing what it is doing.[11]

There are two different kinds of action, which perhaps ought to be separated here. Many have argued that we need to distinguish non-intentional actions – actions done for no reason – from intentional actions. If there is such a distinction then we might imagine that while our intentional actions are likely to prove accessible to us in the

[10] Experimental psychological data tends to confirm this in so far as it finds that subjects judge themselves to be acting before their bodies can be observed to be moving. If they had to observe the movements of their bodies in order to ground their judgements we would expect their judgements to come much later.

[11] This connection is noted by Olsen (1969) and Broadie (1997). Note that someone who denied that our knowledge of our actions was *relatively a priori* might still accept that they must be *relatively transparent*. They might hold that it is a condition of something's being an action that it be generally known to the agent, but hold that our actions are known to us on the basis of perceptual information. They would then have to say that when the subject is blind to their movements, their movements cannot be actions.

way characterized, our non-intentional ones may not.[12] Let us count absent-minded finger tappings as non-intentional actions of mine. Am I epistemologically disassociated with such actions to a degree that makes the claim that I could be totally self-blind with respect to them look plausible? It is clearly true that I can be tapping my fingers without noticing. However, to the extent that it is plausible that there is genuine *agency* in such cases, by which I mean that *I* can be said to be controlling the action, I must normally be able to come to know what I am doing. (And note that it is far from clear that it *is plausible* that we have genuine agency in such cases, rather than just that my body is caused to move by a purposive system I am not responsible for.) Consider a case where a subject is acting – counting using the fingers of her right hand, say – and then acts in a way incompatible with carrying on doing the former – picking up her cup of tea with her right hand, say. If we want to say that the subject herself is *genuinely* controlling both actions, rather than that some sub-personal system is, we want to explain how the second action relates to the former in the case were the subject is supposed to have no possibility of accessing the former. If there were a general disassociation for the subject between the perspective from which the question 'what am I doing?' gets answered, and the perspective from which she carries out her actions, we would have the possibility of a subject deciding to pick up her cup of tea wondering if her action could bring to an end or disrupt, or indeed be disrupted by, other actions she was unaware of carrying out. But turning the subject's actions into possible external impositions in this way seems to get things quite wrong. If what I am doing can be said to be controlled by me, I must at least have the power to initiate it or to will it to cease when I have reason to do so. The control and regulation of my actions as the actions of a unified agent seem to require this. And surely if I have the power to initiate or to stop what I am doing, then what I am doing must normally be in some way accessible to me. Thus for an action to be within a subject's control, the subject must be capable of knowing what they are doing. And given the ability to ask the question 'What am I doing?' the subject's awareness of what they are doing must normally feed into an answer.

[12] See, for example, O'Shaughnessy (1980, Vol 2, Chapter 10) and Ginet (1990, p.3). By a non-intentional or sub-intentional action is meant not just an action which is not intended under one description but an action that is intentional under no description.

With these features of our knowledge of our action identified, I want now to explore some possible accounts we might give of our knowledge of actions and to sketch out the approach I think most promising. Some of the accounts explored do not sit well with the intuitively plausible features identified, and some have other shortcomings. However, nothing I say here about the accounts I set aside, in favour of the one I think most promising, will be sufficient to show that they could not be made good, or could not explain, or explain away, the features identified above. My hope, nevertheless, is that enough critical work will have been done to motivate a consideration of the account I think we should go for.

3. Possible Solutions

3.1. The Dual Component Model

Certain theorists about belief have attempted to solve the problem of how it is that we have relatively a priori knowledge of our beliefs despite their relational character, by claiming that what we took to be a unified phenomenon is in fact a dual component one.[13] Thus our knowledge of our beliefs on this account divides into two parts: knowledge of a narrow component, of which, as subjects, we have a distinctive kind of knowledge – knowledge by introspection, and knowledge of an external component not knowable by the subject in any distinctive way. The dual component theorist tries to explain how it is that our knowledge of our beliefs has the features it has, despite being relational, by adopting a dual component theory of belief and an account of our knowledge of each element. If this is a natural move to make in dealing with our knowledge of our beliefs, it is an even more natural move to make in dealing with our knowledge of bodily action. It has been common to think of actions as the combination of at least two separate components: a psychological component – an intention, or trying; and a non-psychological one – a movement of the body. If this were to be our view of actions then we could say that the features that we have taken to be features of our knowledge of our actions are in fact only features of our knowledge of our tryings or intentions. Our knowledge of our intentions is authoritative, a priori, and transparent, but our knowledge of our bodily movements is not. As Donnellan puts it:

[13] Putnam's original paper, Putnam (1975), suggested this move. It is also developed in McGinn (1982).

> What this suggests is that our knowledge of our own intentional actions is complex, that it divides up, so to speak, into an element of 'direct awareness'... and other elements to which observation is relevant.[14]

To try to explain the knowledge we have of our actions by claiming that it is in fact knowledge of two more primitive components is not, I think, the right direction to move in.

First, such an account makes a critical part of my knowing what I am doing a matter of perception. Given that we have perceptual access to other's bodily movements we might think that it fails properly to capture the first-person/third-person asymmetry involved in the first-person authority we have found to be a feature of our knowledge of our actions. Also, such a dual component account clashes with the seeming aprioricity of our knowledge of our actions. The account has it that my knowledge that, for example, I am raising my arm depends not only upon my knowing immediately that I intended or am trying to raise my arm, but also upon perceptual information that my arm is rising. In so doing it fails to explain a feature of our knowledge of our actions we found plausible. Moreover, it threatens to conflict with the feature of relative transparency. If my knowledge of my actions depends directly upon a perceptual component in this way, then given the possibility for widespread brute error that perceptual capacities leave room for, we should expect to find such a possibility in the case of action. It maybe that the dual component theorist can explain why we must normally be able to know what we are doing despite adopting a partially perceptual model. They may argue that since our normal conditions for initiating action require information about the nature of our bodies, via perceptual feedback, perceptual breakdown would in fact rob me of my capacity to act, by robbing me of a grasp of what possible things I could do. If this were the case, there would not be the possibility of actions to which I was self-blind in the case of perceptual breakdown. However, even given perceptual breakdown, we may think that a subject would in fact be able to continue acting for a short time, relying on past information for a grasp of the possible actions it could carry out. And it is not plausible that they would be unknowing about such actions simply in virtue of that breakdown. The theorist might also have some other reason for saying that there is a conceptual dependence between our acting and our knowing that we are acting, so that there would be no actions in the

[14] Donnellan (1963).

case of perceptual breakdown. *Prima facie,* however, it is hard to see what their reason might be. But the reason cannot be the attractive one that there is something about the conditions under which we act which is already apt to provide us with knowledge of the action. This cannot be their reason because, on their account, to know what I am doing I have to wait to receive perceptual information which is not available at the onset of the action about what my body is doing.

Second, the dual component model seems to be partly motivated by the assumption that an account of the knowledge we have of our intentions or tryings, an account which was supposed to explain the distinctive features of action self-ascription, can non-problematically be provided. The assumption seems to be that our intentions or tryings will involve no external component and will be capable of being the objects of direct awareness. However, first, we would need to make sense of such direct awareness of any supposed internal components of our actions, and, second, we have no good reason to think that the contents of our intentions or tryings will be any more independent of their relations to external elements of the environment than the contents of our beliefs and perceptions. It is highly plausible to think that a creature inhabiting an arm-free world cannot even try to raise an arm.

Third, the model can seem unsatisfactory for much the same reason that a dual component theory of belief, and of the knowledge we have of our beliefs, has struck theorists who think of belief as a unified psychological state as unsatisfactory. If we think of actions as unified psychological phenomena, rather than aggregative constructions of psychological and non-psychological components, then we are not going to adopt the view outlined above. As I have said, I am not going to offer a defence of this view of actions here but merely consider what account of our knowledge of them will seem plausible for someone who accepts such a view.

Considering how a dual component theorist will respond to the problems raised, brings out the fact that are in fact many different forms of dual component account. There are as many different forms as there are combinations of the different accounts one could give of the nature and knowledge of a person's knowledge of their intentions or tryings, on the one hand, and of nature and knowledge of the movements

of their body, on the other. The dual component model envisaged above is only one version of a dual component account.

.

In dealing with the first problem raised above, it is important to stress that a contemporary dual component account is likely to hold that our knowledge of our actions depends not just on our five senses, but also centrally on bodily awareness.[15] And given a view on which bodily awareness is necessarily awareness only of my body, there is room to say that, even in the perceptual component, there is a first/third-person asymmetry. So, there is room to say that first person authority can be reflected in both components of the dual component account. However, we still have a problem with relative aprioricity – it will still be the case that that I need to await the deliverance of my perceptual capacities in order to know whether I have raised my arm. This has seemed to be phenomenologically counter-intuitive. It rather seems that my raising my arm, say, gives me *ceteris parabis* the knowledge that I moved my arm without my having to monitor feedback from bodily awareness. It has not been doubted that this knowledge is set against a background which makes it unnecessary for me to check to see whether my motor system is in working order. Or that having reason to think that our motor system is playing up is very likely to result in an appeal to the testimony of our senses to check whether we acted. Nor is it doubted that keeping up an accurate schema of what actions are possible, and so being able to continue acting at all, will require perceptual feedback on the position of my limbs and so on. However none of this gives us reason to think that any such feedback is part of my immediate grounds for my knowledge that I am raising my arm.

We can imagine our dual component theorist, in response to the second problem raised above, rejecting the assumption that our knowledge of our intentions or tryings will involve direct awareness of some component, narrowly construed. What account the dual component theorist comes to offer of our knowledge of our intentions or tryings will depend on what account she thinks right for mental phenomena in general. So, instead of taking our knowledge of our mental states and actions to be due to some kind of internal perception, she might think it due to some reliable

[15] Of course we can also come to have knowledge of the movements of our bodies from observing and listening to others – someone starting can tell me that I have touched them, or someone can verbally inform me that I have moved – but such sources are clearly not operative in the basic case.

mechanism. Or she might think that our intentions or tryings can function as reasons for their own ascription. Whatever account is adopted (and it may be an account that embraces an externalist individuation of intentions and tryings), it may then be combined with an account of our knowledge of our bodily movements, and perhaps our knowledge of the relation between our bodily movements and our intentions or tryings, to give an account of our actions. However, it needs to be noted that once the theorist has embraced the idea that intentions or tryings depend upon external elements, much of the motivation for giving an account of our knowledge of our actions by thinking of actions as divided into two components falls away.

However, recalling the third problem, even given adaptations of the kind discussed, if one operates with an understanding of actions as unified phenomena, then however compelling in themselves the component parts of the account offered are – the account of our knowledge of our intentions, of our bodily movements and so on – the resulting account of action is unlikely to be accepted.

So far we have supposed that our dual component theorist is committed to a dual component theory of action and, thus, to a dual component account of our knowledge of our actions. Might it not be possible for a theorist to meet the third problem by agreeing that our actions are unified psychological phenomena, but holding that our knowledge of them nevertheless comes in two parts? Here the suggestion would be that I have knowledge of my intentions or tryings, however such knowledge is construed, and also knowledge of my bodily movements, through bodily awareness and/or ordinary perception, and that these two kinds of knowledge are what enable me to know what I am doing. On this account my knowledge of what I am doing will be essentially indirect. According to this version of the dual component account, intending or trying is at most a necessary condition of, but not a component part of, an action. So, on this account my knowledge of my intending or trying is knowledge only of a necessary condition – perhaps of a causal antecedent of my action – not of my acting. And any observation of my body moving will not in itself give me knowledge of what is in fact my action. If we held a unified account of the nature of bodily action, but adopted a dual component account of our knowledge of such actions, we would in effect adopt an account on which our knowledge was always

inferred from the knowledge that necessary conditions for my acting were satisfied. Given the continued presence of a perceptual condition such an account would again conflict with the relatively a priori nature of our knowledge of our actions claimed above, and threaten to conflict with the transparency claim. However, identifying this possibility does raise the question of whether our knowledge of our actions is some kind of inference from knowledge of necessary conditions, combined with background conditions.

3.2. The Inference Model

The suggestion that we appeal to an inference to explain our knowledge of our actions can only be the suggestion that we make inferences about our actions from either our observations of our bodies or from our knowledge of the precursors of actions: intentions or tryings.

The idea in the former case would have to be that I come to know what I am doing from making an inference on the basis of perceptual information, through bodily awareness and the other senses, about the movements of my body. It seems clear that, if the remarks made above against the dual component theorist's idea that part of our knowledge of our own actions flows from our sensory access to our bodily movements were well taken, then this account will also be held in doubt. It will sit well with neither the seeming aprioricity of our knowledge of our actions, nor with their relative transparency.

What of the suggestion that we infer that we are acting from our knowledge that we are intending to act, or trying to act, together with an assumption that our motor systems and so on are all in working order? Well, first, it is very important to stress that trying, in this context, has to be understood, much like an intention, as an independent precursor of my actually acting. If the picture is that my trying is *part* of my action, then it would conflict with our unity thesis for actions. If the supposition is that my trying *is* my acting then, whilst there may be objections to the identification, the attempt to give an account of the trying would be to give an account of the action. Given this clarification, on this version of the inference model I have knowledge of my actions via direct knowledge only of my intentions or tryings and the assumption, say, that my motor system is in working order. Consider the account as told for

intentions (and given the idea of my trying as an independent precursor of my acting, I think that what will be true for intentions will be true for tryings).[16] On this account, I know that I am doing φ indirectly as an inference from knowing that I intend to φ, and that my motor system is in working order. The first thing that strikes us about such an account is that it does not fit our naïve understanding of what is going on when we know we are acting. It does not *seem* to me that I know I am raising my arm as a result of inferring that I am from my knowing that I intend to raise my arm; it rather seems that I know that as a result of my raising my arm. On this account, no further epistemic support is given by my doing anything. Further, it seems impluasible that the subject need have beliefs about their motor system in order to be credited with knowledge of what they are doing. But even putting aside these considerations, without more being said, the grounds given in the account do not seem to be sufficient to ground knowledge that I am φ-ing. Maybe I usually do what I intend to do, but all too often, even if I could do what I intend to do, I do not. It would seem that my intention to φ, plus assumptions about my motor system working properly, may plausibly ground a justified hypothesis that I will φ, but it does not seem able to ground my *knowledge*. The above inference models do not seem easy to accept. There are more sophisticated variants of the latter form of the inference model. These, however, fall under what I call the anticipation model which will be discussed later on.

3.3.The No-Reasons Judgemental Model[17]

In reviewing accounts of self-knowledge with respect to belief, Peacocke identifies what he terms 'no-reasons accounts'. According to no-reasons accounts, our beliefs about our beliefs count as knowledge although they are held on the basis of no reasons. In particular, the first-rder belief does not constitute, on such an account, my reason for my self-ascribing my belief.

The simplest of such accounts has it that as a matter of fact agents like us, due to some internal mechanism, only self-ascribe beliefs that they actually have. Given that

[16] If my trying to do something is not supposed identical with, or part of, my acting when successful, then in order to preserve the assumption that my acting is my doing something we would have to assume that trying is not all I have to do to act. But if we assume this, then it seems that the gap between tryings and actions will be much the same as that between intentions and actions.
[17] I take this description from Peacocke (1998) and (1999).

a self-ascription of the belief that **p** is a reliable indicator of a belief that **p**, if we accept reliablism about knowledge, it can be said that by such self-ascriptions agents like us express knowledge of their beliefs. The equivalent account for the self-ascriptions of actions would have it that as a matter of fact we tend be reliable about what we are doing, due to some non-epistemic feedback mechanism, and in virtue of that we are capable of knowing what we are doing.

Such simple reliablist accounts are acceptable to few in making the connection between the psychological phenomenon, and the knowing of it, an entirely brute connection. Such accounts tend, therefore, to be embellished in such a way that the relation between the phenomenon and the knowing of it is a rational relation, either because it is partly constitutive of the relata, or because it is taken to have some wider rational significance. Shoemaker argues that it is not just a matter of fact that belief ascriptions are authoritative, and beliefs relatively transparent, it is rather a conceptual, and so necessary, truth that they are. It is such a conceptual truth because it is partly constitutive of what it is to be a belief that a creature with the belief and capable of self-ascription, will normally be able to self-ascribe it. Burge argues that what makes the connection between our beliefs and our self-ascriptions capable of sustaining knowledge is the role such a connection plays in our nature as rational beings. He argues that our nature as critically rational subjects depends upon the obtaining of such a relation and takes this fact to confer a kind of external warrant for the self-ascription.[18]

Similar embellishments could clearly be proposed for our account of our knowledge of our actions. We might take actions to be self-intimating as a matter of conceptual fact. So movements of our bodies that we were normally unable to self-ascribe would not count as actions. We might also claim that an ability to be authoritative about its own actions is a prerequisite for a critical agent, and take this to be sufficient to confer the status of knowledge on to the agent's self-ascriptions.

Whilst both suggestions have considerable plausibility about our natures as agents, the no-reasons theorist faces the same question in this case as he faces in the case of

[18] See Burge (1996).

belief and that is: how do these truths secure *knowledge* for the subject when they might seem to suppose it?

With respect to the first suggestion, we want to say that, while it may be true that certain processes or activities only count as actions if they are epistemically accessible to the subject in a certain way, this does not mitigate the need for an independent account of the way in which they are epistemically so accessible. With respect to the second, it is hard to suppose that the self-ascription counts as knowledge merely in virtue of the fact that we need to take it as a knowledge claim in order to consider the self-ascribing subject as a critical agent. The essential role played by the self-ascription within the critical agent may confer some warrant on it. However, we would expect some epistemic relation to hold between the action and its self-ascription independently of that role, and expect it to explain, at least partly, why the self-ascription counts as knowledgeable. That explanatory task still seems pressing.[19]

3.4. The Anticipation Model

One obvious way to make true our beliefs about an area of which we have control, is to bring about what we believe rather than have our beliefs conform to the facts as they independently are. Our actions are, paradigmatically, an area over which we have control. One suggestion for why we seem to be authoritative about what we are doing, and why we are able to know what we are doing without any process of evidence gathering, is that we do what we believe we are going to do. Velleman, in his *Practical Reflection*, advances a sophisticated development of this suggestion. Velleman advances the basic thesis: that our knowledge of our own actions has the features it has because we normally do what we believe we are going to do. However, his account also aims to meet the question as to why we should be motivated to do such a thing – why do we not just wait and see what we do and form our beliefs about what we are doing accordingly? The answer is taken to lie in a deeply rooted desire for self-knowledge, in particular a desire to know what we are doing. Given a desire to keep up with what one is doing we will be motivated to bring it about that we accurately foresee what we are going to do, and that motivates us actually to do what

[19] These are in effect Peacocke's criticisms of Shoemaker and Burge. See Peacocke (1999, Chapter 5). Despite this worry about Burge's account, it will be obvious that much else that is said in this paper is influenced by the discussion in his (1996), and also in his (1998).

we foresee doing. Furthermore, Velleman identifies intentions with self-fulfilling expectations: to form an intention is to form a belief about what I will do, that I am poised to make true on account of my desire to know what I am doing. Given a framework in which I do what I expect to do, a most effective way for me to get myself to do something – which is after all what an intention is – is to expect that I will do it.

Velleman's account has the attractive feature that it explains the transparency of our actions and the spontaneous, relatively a priori, nature of the knowledge we seem to have of them, while also allowing that the authority we have over our self-ascriptions is not groundless. He distinguishes between having knowledge that is adduced from evidence and knowledge which is supported by, though not adduced from, evidence. We are, on his account, justified in claiming to know what we are doing, not because we form the belief about what we are doing on the basis of evidence, but because we form the belief within a framework which we know makes the belief true. Self-knowledge is a kind of justified invention.

Before considering this account further, it is worth noting that it too has a parallel account in the literature on the self-ascription of belief. Accounts of belief self-ascription that have it that the reason our self-ascriptions tend to be authoritative, while appearing spontaneous or groundless, is that what in part determines the beliefs we have is what we are inclined to self-ascribe. If our beliefs are in this way ascription dependent – we believe what we take ourselves to believe – then there is no epistemic gap between our belief that **p** and our taking ourselves to believe **p** that needs to be bridged by reasons or evidence.[20] We need to be careful, however, to distinguish three distinct claims that may be conflated. First, few would doubt that our self-ascriptions are determinative of our beliefs merely in the sense that answering the question: 'Do I believe **p**?' will tell me, almost always, whether indeed I do believe **p**. This is just first-person authority. Second, few would doubt that settling an answer to the question: 'Do I believe **p**?' will not only tell me whether I do indeed believe **p** already, but will also settle any questions I may have of whether to believe **p**. As Evans has made clear, our procedures for determining whether we believe **p** are the

[20]See Wright (1989) and (1998).

same as the procedures for determining whether **p**.[21] However, neither this claim, nor the former one, commits us to the view that my second-order beliefs, my beliefs about what I believe, are in any way epistemically or ontologically prior to my first-order beliefs. Rather, what they make clear is that our beliefs, and our beliefs about our beliefs, are in a kind of concord. Third, and more contentiously, in claiming that our belief ascriptions are determinative of our beliefs we might be saying, not just that the same procedures can be used to answer both questions, but rather that the process of determining what I believe is, in the normal case, what brings my beliefs about. This makes my self-ascriptions ontologically and epistemologically prior to my beliefs, and introduces a kind of psychological idealism in being committed to there being no fact of the matter about what I believe prior to my reflections explicitly about what I believe.[22]

Velleman's idea that our knowledge of our actions is somehow invented has most in common with the third claim identified: in both cases our self-ascriptions have the features they do because the self-ascriptions make the facts fit them. However, there are of course differences between actions and beliefs which make a significant difference to the acceptability of the accounts being offered. While actions are single occurrences that can be willed, beliefs are states which cannot. Thus, while we can make sense of a system that brings it about that it does φ, because it anticipates doing φ and desires self-knowledge, it is very much harder to make sense of a system that brings it about that it believes **p** somehow *because* it self-ascribes the belief **p**. The latter suggestion makes our self-ascriptions appear ungrounded and our beliefs, therefore, either arbitrary or a matter of will.

However, even if not problematically idealist, the trouble with accounts of the sort Velleman proposes, is that they are – as he fully realises – bound to seem to have things backwards. We seem to have reasons to do things that are not grounded in beliefs about what we take our selves to be about to do. Indeed, we could have

[21] Evans (1982,p.225)*:* 'I get myself in a position to answer the question whether I believe that **p** by putting into operation whatever procedures I have for answering the questions whether **p**.'
[22] An account such as this can easily give way to a kind of psychological instrumentalism – rather than taking it that there really are beliefs constituted by the practices of self-ascription, we might think that there are only the practices of self-ascription themselves, and that the beliefs ascribed are useful interpretational fictions.

reasons to do things without any capacity for first-person higher-order beliefs. Further, as we noted before, looked at naïvely our beliefs about what we are doing are given epistemic support by what we are in fact doing. Velleman, as we have seen, is keen to emphasize that an account of the kind he offers does not need to hold that our self-ascriptions are epistemically groundless and lacking in justification. Nevertheless, the justification a subject has for their belief that they are doing φ cannot lie in the fact of their doing φ. Rather, it must lie in the fact that the belief is held in circumstances in which, given the subject's desires for self-knowledge, the belief will ensure that they are motivated to do φ. While it seems right to say that such nesting of a belief can provide evidential support for the belief, it is hard to accept that this is the way epistemic support figures in the standard case.

Another important cause for concern is that this account works only by assuming that intentions, understood as beliefs about what we are going to do, are necessary for action. That makes the account unusable to someone who thinks we need form no such belief about our future action as a result of practical reason. If we hold the position mooted before, that we can sometimes just act as the conclusion of a process of deciding how to act, then we are not going to want to hold that a belief about what we are going to do is necessary for our knowledge of that action.

What an account such as that offered by Velleman does seem to get right, however, is that in cases of psychological self-knowledge, the knowledge we have of the phenomenon is somehow simultaneously given with its occurrence, rather than merely occasioned by it. As such, it seems to respect the features of our knowledge of our action intuitively identified. If we cannot reconcile this fact with the naïve thought that we know what we are doing partly because we are doing it, then we may be forced to adopt such an account.

5. Knowledge Through Participation

In what follows I will attempt to provide a sketch of what a more plausible account might look like. Let us go back to the naïve thought that our beliefs about what we are doing are given epistemic support by what we are in fact doing: we know what we are doing *because* we are doing it. Now let us suppose given the arguments presented

above, that our action cannot function in our justification for its self-ascription in virtue of some observation of it, or inference from its effects or causal antecedents. What of the possibility that the action can function unmediated as the reason for its own ascription?

Let us then consider a view according to which we know we are acting when we are, because, and for the reason that we are.[23]

Now presented merely as such, the account faces a problem. It provides us with no explanation of how an action, in so far as it need only be world directed, can bear any normative relation to our judgements about ourselves to the effect that *we* are acting. To put it simply, what is it about Smith's raising his arm that can support Smith's knowledge that he is raising his arm?

Well, we may say at this point that it is not simply Smith's raising his arm that can function as the reason for Smith's self-ascription of the fact that he is raising his arm. Rather, it is centrally those actions Smith is consciously engaged in which stand to function as the reasons for their own ascription. The crucial question then becomes: how should we understand the sense in which the action is one that Smith is consciously engaged in? We might try two ways: (i) raising an arm, say, might seem to give us consciousness of the arm. This does not seem to help. (ii) We might take a conscious action to be an action that we make a judgement about. But, then we have gone a full circle in our explanations.

Let us again pick up the suggestion of Peacocke's about how to characterize what is involved in the consciousness of a conscious attitude. The key suggestion made by Peacocke is that a state is conscious, in the relevant sense, if it *occupies our attention*. He tells us that to understand properly what is meant by a state occupying our attention, it is important that we distinguish it from the case of a state being the *object of our attention.*

[23] The suggestion is, of course, drawn from Peacocke's (1999) account of our knowledge of our beliefs.

So, how are we to understand what it is for a state to occupy, without being the object of, our attention? Peacocke's discussion at this point does not offer a further explication. He identifies a number of distinct mental phenomena that can be said to occupy without being objects of attention – perceptions, conscious attitudes, tryings and actions – but does not attempt any further account.

However, we clearly need more of an account, and such an account may prove to be easier to come by in the case of actions than in the case of Peacocke's conscious attitudes.[24] The things that we do strike us immediately as the paradigm of things that occupy, without being the objects of, our attention. They are the paradigms of the things known by attentive engagement or participation, rather than because they function as the objects of some scrutiny.

How then might we understand what it is for an action of mine to occupy my attention? That is, how should we understand what is going on when we say that an action is conscious in virtue of the subject's engagement with it? I am going to assume that dynamic entities such as actions and events have properties, and that an actions being conscious is a matter of the action having a certain property. I am also going to assume that we can make a distinction between the relational and intrinsic properties of an action or event. So, to take an unproblematically physical example, we would I think want to say of the rolling of the pink ball into the pocket of the table, that its 'being a fast rolling' is a relatively intrinsic property of the event. On the other hand, its 'being a rolling which occurred after the sipping of the beer' would seem clearly to be a relational property of the event. It is of course hard to state any principled difference between relational and intrinsic properties (especially since it is likely to be a partly contextual matter whether a property is considered intrinsic or not). But roughly, in order to fix the intrinsic properties of an event, say, we need consider only the event itself along with standing background conditions, and can bracket other independent events that may be occurring. What this suggests is that, for our purposes, we can think of a property as an intrinsic property of an action if it is a genuine monadic property of the action (even if requiring relational individuation).

[24] In fact I think that what Peacocke calls a conscious belief is in fact a kind of mental action; and that the account required of what enables our beliefs to function as reason for their own ascriptions will be

And we can think of a property as a relational property of an action if it stands to be analysed as a dyadic property, that is, as a relation.

Keeping in mind this distinction between the intrinsic properties and the relational properties of a dynamic entity such as an action, and the assumption that being conscious is a property of an action in virtue of which it occupies our attention, we can ask whether such consciousness is to be understood as a relational or as an intrinsic property of the action. What answer we give to the question depends upon the model we have for what it takes for an action to occupy our attention. Let me consider two possible models.

We might think that when we engage in an action in such a way that the action occupies our attention, what we have is a kind of complex action involving two more basic actions. First we have the action of doing φ and second we have the action of attending to doing φ. So our doing φ in a way that makes it occupy our attention is a matter of its having the relational property of being part of a complex action which involves not just a doing of φ but also an action of attending to a doing of φ. Here doing something consciously is a matter of its standing in a certain relation to something else I do. To adopt this as an account of what is involved leaves us, of course requiring an account of what is involved in the subject attending to doing φ where φ *occupies* attention rather than being the object of it.

We might think that when we do something in a way which occupies our attention what we have is a single action done in a certain way. Rather than there being two actions, the doing of φ and the attending to the doing of φ there is rather just the one action; the doing of φ in a certain way. Here the suggestion is that for an action to be conscious in the sense that it can stand as the reason for its own ascription, is for it to have an intrinsic property.

It is very natural to assume that for a mental state or activity to be conscious in a way that makes it accessible to its subject – for it to be conscious in a way that makes it

---

parasitic on an account of what enables our judgements to function as the reason for their own ascriptions.

poised to stand as the reason for its own ascription – is for it to have a certain relational property, for it to bear a relation to some further act on the part of the subject. In particular, it is thought that it can only have such a property in virtue of some further act of attending to their actions, for example, on the part of the subject. This very natural thought is, in my view, a mistake (at least as far as our *activities*, as opposed to our mental *states*, go).[25]

Instead, I want to suggest, we should take the kind of consciousness exhibited by mental activities, which occupy without being object of attention and are able to stand as reasons for their own ascription, as intrinsic to those activities and not requiring any further mental act on the part of the subject. The kind of consciousness inherent in my actions should be thought of as a monadic property of the action rather than a dyadic one.

However, with only this much in place we are still left with two questions:

What *way* of acting makes an action conscious, and so occupy rather than stand as the object of our attention?

Why should an action that occupies our attention qualify for knowledgeable *self*-ascription? What grounds my taking the action I am consciously carrying out to be *my* action?

What seems clear is that an adequate answer to the first question should deliver an answer to the second question. If an action's being conscious makes it fit to stand as the reason for its own self-ascription, then what it is for the action to be conscious should make it clear why it does ground such self-ascriptions. In essence, my suggestion will be that I act consciously when I engage in my action as something I control. Further, that engaging in an action as something I control is engaging with the action as *my* action, and so involves a primitive form of self-awareness. Because of this my conscious actions are apt immediately to ground self-ascriptions.

[25] It is in fact my view that passive mental states such as beliefs and perceptions are only conscious in this sense (as opposed to being conscious in the sense of making the subject conscious of objects in the

Actions strike us as the paradigms of mental phenomena that can occupy or engage our attention without being an object of it. Quite intuitively, our actions are those things we know, not by observing them, or by reflecting about them, but rather by actively engaging in them. We, of course, carry out many actions as a matter of habit, or as relatively automatically and unattentively, but when we act consciously we seem to act with a sense of guiding our action, with a sense of control. A natural suggestion as to what distinguishes an action carried out in a way that occupies our attention, from an action carried out in a way that does not occupy our attention, is that an action which occupies our attention is one the agent carries out with a sense of control.

But what is it for an agent to act with a sense of control? This is obviously a large question, but to act with a sense of control must, at least be, to be aware of guiding our actions, to experience our actions as those we initiate and those we have the power to stop. We can either take this kind of awareness as primitive, or attempt further explication. If such explication of what it is to act with a sense of control is available it will, I think, most likely rest on our acting with an awareness of the means by which we control or guide our actions. In essence, we control our actions by acting on the basis of an evaluation of the possibilities open to us. So, one suggestion would be that we experience our actions as controlled when we act on the basis of our evaluation of possible actions, grasped as possible. On this suggestion, to experience an action as controlled is for the action to be the result of a process of evaluation of the options available, grasped as options. Conscious bodily action would then involve the agent having a grasp of the possible ways that they could move their body as basic actions, and carrying out one action rather than another on the basis of an assessment of their options. The agent's grasp of which actions are available as basic actions will be based on a general grasp of the ways in which they can move their body, which itself will be based in ways that they have moved it in the past. It will also be based on a particular grasp of the position of their body at the time of action, which itself will be based on the ways in which they have most recently moved it. Let us take a case of an agent with only two options: suppose that an agent has just one barely functioning arm that they can move in just one way. If the agent grasps the two options of 'moving

---

world) because they bear certain relations to non-passive states such as judgings and perceivings.

arm up' and 'not moving arm up' as things that could be done, and acts directly on the basis of an assessment of these options as 'to be done' or 'not to be done', the subject seems to have what is needed to act with a sense of control. However, if this is what is involved in acting with a sense of control, then we have reason to think that acting with a sense of control gives us some awareness of my action *as* my action. It is certainly the case that only the agent of an action can be engaged in their action in this way. Only the agent of action can act immediately on the basis of an assessment of options. But we also have reason to attribute to the agent awareness not only of the action through this kind of engagement with it, but also an awareness that it is *their* action. Acting directly on the basis of an assessment of options, grasped as options, means I have grasped that action as an option for *me*. An agent may not be able to *self-ascribe* something as an option, because it may not have the requisite first-person concept. However, acting directly on the basis of an assessment of whether to carry out an action or not, given a grasp of the choices, must either presuppose, or give me, a grasp of the action as an option for me.

We can summarize the line of thought using the following argument:

1. An agent acts with a sense of control when they carry out their action on the basis of an assessment of the options, grasped as options, of acting one way rather than another.
2. If an agent acts directly on an assessment of the options, grasped as options, of acting one way rather than another, they manifest awareness of the option as an option *for them*.
3. Therefore, acting with a sense of control is acting with self-awareness.

There is of course much more to be said about how to understand these various claims. In particular, it would I think emerge from a proper consideration of them that there will be various kinds of grasp that can be attributed to a subject of their options, and that there will be correlated notions of self-awareness. There will be agents who can fully conceptualize what actions are possible as basic actions, and who deliberate over their reasons for acting one way rather than another. And there will be agents who can have only a more primitive grasp of their options and who evaluate them in less sophisticated terms. A creature who manifestly experiments and revises their

strategies for getting what they want, but who cannot be thought to fully conceptualize their options might fall here. There will also be agents who function as rational systems in such a way that it seems right to say that they in some sense act on the basis of an assessment of their options, but for whom we do not want to say that they grasp their options *as options*, and so who do not act with self-awareness.

If the above suggestion has anything going for it, we have the beginnings of an account of our knowledge of our own actions. We are able to act consciously, and acting consciously is acting with a sense of control. Acting with a sense of control is acting in a certain way: it is acting directly on the basis of an evaluation of the possible ways of acting, understood as possible actions. Acting directly on the basis of an evaluation of the possible ways of acting, understood as possible actions, is acting with self-awareness of a primitive form. So, when my acting consciously acts as the reason for my self-ascription of the action, my self-ascription is knowledgeable because it rests on an awareness of what I am doing.

I have said that acting consciously, which is acting in a certain way which grounds our self-ascriptions, involves a form of self-awareness. This may give rise to objections. It might be said that properly speaking such consciousness cannot be a form of *self*-awareness as it does not involve a capacity for first person reference.[26] Certainly if being self-aware is understood to imply a capacity for first person reference then this cannot be a form of self-awareness. However, it is form of awareness which is such that a suitably cognitively equipped subject – a subject with grasp of the first person and the concept of an action – will immediately be able to *self*-ascribe the action they are conscious of in this way. A form of awareness which is self-indicting in this way clearly needs to be distinguished both from our awareness of the world and things around us, an awareness that may in many cases be quite independent of our ability to ascribe anything to our selves, and from our self-ascriptions themselves. I do not mind whether we call this primitive self-awareness, or something else. The point is the need to identify the need for and nature of the phenomenon.

[26] This is what for example Casteneda would say.

By way of conclusion, let us see how the account sketched fits with the features of authority, relative aprioricity and transparency introduced at the beginning of the chapter as intuitive marks of our knowledge of our actions.

Given the claim that, in central cases, it is the acting in a certain way that grounds ones knowledge of ones actions, we have an explanation of first-person authority. The agent whose conscious or engaged action stands as the reason for her self-ascription will be first-person authoritative over her actions. The agent of an action will know in a way unavailable to others whether she is acting, because only the agent acts directly as a result of her assessment of the possibilities available, understood as possibilities. Furthermore, the agent will not only in general know that she is acting when she knows in this way that she is acting, she will also know what she is doing. As long as she has no evidence to the contrary, the agent is entitled to the assumption that her motor systems are working properly and that she has, on the basis of past action, a veridical grasp of both the general and particular possible basic actions open to her. In bringing about a movement of her body as the direct result of an evaluation of the ways she might have moved it, grasped as ways she might move it, the agent can be said to know what she is doing. She can know what she has done in a way not dissimilar to the way she could know which object she has picked out if she had a grasp of the possible objects available and picked one, rather than another. However, it is however, important to emphasize a couple of points made earlier. When it is claimed that our actions can act as the grounds for our knowledge of them, it is not being claimed that there is knowledge of them under any description. So, while the moving of my hand may be the potting of the pink ball, my knowledge that I am so moving it, does not by itself give me knowledge that I am potting the pink. The claim of first-person authority with respect to our actions is to be understood as relative to certain descriptions which could be regarded as basic. Authority is also compatible with the possibility of mistaken basic action ascriptions, of raising my arm, for example. My motor system could malfunction in such a way that my arm does not move, or that it lowers rather than rises, or I could for one reason or another have a non-veridical grasp of the possible actions open to me as basic.

It should by now be clear how our knowledge of our actions is being taken to be compatible with the feature of relative aprioricity. There is obviously a considerable

role that is played by our proprioceptive and other perceptual faculties in maintaining and updating a subject's grasp of the possible ways she can act. However, when an agent with a grasp of the possible actions available carries out a single basic action, she need not avert to the testimony of her senses to know what she is doing. While our perceptual faculties are clearly required to give me knowledge of the things I might do, they are not required to give me knowledge of which, out of the things I might do, I am doing. All that is required to give me knowledge of that, given the appropriate background, is to do it.

What of relative transparency? Given that acting consciously has been understood to be acting on the basis of an evaluation of one's options, grasped as options, we have reason to think that any agent who acts while asking themselves what they are doing will act consciously. An agent who acts on the basis of an assessment of their options while considering what they are doing, will be an agent who is asking what options for action are being taken while assessing those options. That must mean that the agent acts on the basis of an assessment of their options, understood as options. Given that acting consciously will, in the absence of any repressive mechanisms, be sufficient to ground a knowledgeable self-ascription, an agent who acts, and who asks themselves what they are doing, will know what they are doing.

Bibliography

Anscombe, G. E.M. (1963), 'Two kinds of error in action', *Journal of Philosophy* 60:393-401.

Broadie, F. (1967), 'Knowing that I am doing', *The Philosophical Quarterly* 18:137-149.

Burge, T, (1996) 'Our entitlement to self-knowledge', *Proceedings of the Aristotelian Society* 96: 91-116.

Burge, T, (1998) 'Reason and the first person' in MacDonald et al (1998).

Cassam, Q. (ed.) (1994), *Self-Knowledge*, Oxford: Oxford University Press.

Castaneda, H-N. (1999) *The Phenomeno-Logic of the I: Essays on Self-Consciousness*. Bloomington and Indianapolis: Indian University Press.

Danto, A.C. (1963), 'What we can do', *Journal of Philosophy* 60:435-446.

Donnellan, K. (1963), 'Knowing what I am doing' *Journal of Philosophy* 60:401-409.

Dunn, R.(1998), 'Knowing What I'm about to do without evidence', *International Journal of Philosophical Studies* 6 (2): 231-252.

Evans, G. (1982), *Varieties of Reference*. Oxford: Clarendon Press

Ginet, C. (1990), *On Action.* Cambridge: Cambridge University Press.

MacDonald, C., Smith B. and Wright C. (eds.), (1998), *Knowing Our Own Minds: Essays of Self-Knowledge*. Oxford: Oxford University Press.

McGinn, C. (1982), 'The structure of content' in *Thought and Object* (ed.) Woodfield, A. Oxford: Clarendon Press

Olsen, C. (1969), 'Knowledge of our intentional actions', *The Philosophical Quarterly*; *XIX*, p.324-336

Roessler, J. (Chapter 18, this volume) 'Intention and the First Person'.

O'Shaughnessy, B. (1963), 'Observation and the will', *Journal of Philosophy* LX

O'Shaughnessy, B., (1980) *The Will*: *A Dual Aspect Theory.* Cambridge: Cambridge University Press.

Putnam, H. (1975), The Meaning of "Meaning"' in his *Mind Language and Reality, vol. II*, Cambridge: Cambridge University Press.

Peacocke, C. (1999), *Being Known*, Oxford: Oxford University Press.

Peacocke, C. (1998), Conscious Attitudes, Attention and Self-Knowledge' in MacDonald et al. (1998).

Shoemaker, S.(1996), *The First Person Perspective and Other Essays,* Cambridge: Cambridge University Press.

Velleman, J. D. (1989), *Practical Reflection.* Princeton: Princeton University Press

Wright, C. (1989) 'Wittgenstein's Later Philosophy of Mind: Sensations, Privacy and Intention', *Journal of Philosophy* 26: 622-34.

Wright, C.(1998), 'Self-Knowledge: The Wittgensteinian Legacy', in MacDonald et al. (1998).