

Author Posting. © *Philosophical Explorations* 2011.

This is the author's version of the work. It is posted here by permission of *Philosophical Explorations* for personal use, not for redistribution.

The definitive version was published in *Philosophical Explorations*, Volume 14 Issue 2, June 2011.

doi:10.1080/13869795.2011.569749 (<http://dx.doi.org/10.1080/13869795.2011.569749>)

Perception and Practical Knowledge

John Schwenkler*

Department of Philosophy, Mount St. Mary's University, Emmitsburg, MD, USA

According to G.E.M. Anscombe, an agent's knowledge of his own intentional actions differs from his knowledge of his unintended behaviors as well as the knowledge others can have of what he intentionally does, in being known "without observation". I begin by posing a problem for any conception of this theory according to which non-observational knowledge must be independent of sense-perception, and criticize several recent attempts to get around the problem. Having done this, I develop an alternative account of non-observational knowledge according to which it consists in the particular causal role of an agent's self-awareness in bringing his intentional actions about.

Keywords: Action; Knowledge of action; Self-knowledge; Non-observational knowledge; Anscombe, G.E.M.

1 Introduction

Philosophical thinking about intentional action can tend to pull us toward two independently intuitive but mutually incompatible positions. First, suppose we begin from the idea that the knowledge a person has of his own intentional actions is essentially different from

* Email: schwenkler@msmary.edu

his knowledge both of the actions of others and of any actions of his own that are involuntary or unintentional. Then we can quickly be led to think that the operation of intentional agency is fundamentally an “interior” affair, consisting in some domain of events that a person is guaranteed to know about when they are a part of his own life in ways that he cannot draw on in coming to know about things that happen “outside” of him. Second, however, suppose the action theorist takes his departure from the idea that, as Elizabeth Anscombe puts it, in acting intentionally “I do what happens” (2000: 52), i.e. that human actions are events every bit as “worldly” as the rolling of a stone and the fall of a sparrow. Given this starting-point, a natural place to end up is with the conclusion that one’s knowledge of his own intentional actions can’t be that different from his knowledge of other sorts of events after all, or at least that the epistemic specialness indicated earlier isn’t characteristic of the entirety of one’s intentional behaviors, as opposed to some proper subset of them, perhaps the domain of deliberate bodily movements that one is aware of through “introspection”. Clearly these conclusions can’t both be right, yet each starting point highlights an aspect of intentional agency that seems entirely beyond dispute: is it possible, then, to occupy a position that acknowledges the force of both initial claims but rejects the corresponding conclusions?

It is helpful to think about Anscombe’s *Intention* as, in part, an attempt to show the way out of this sort of dilemma. For Anscombe is unwavering in her insistence that a person’s knowledge of his own intentional actions is distinctive, both in being “non-observational” and in being a species of “practical knowledge”, knowledge that is somehow “the cause of what it understands”. Yet she also rejects any attempt to reduce intentional agency to, or even treat it in a way that privileges our claim to have a special knowledge of, (putatively) “interior” events like intendings or tryings, or even to “mere” bodily movements. If Anscombe is right,

then we can respond to the dilemma sketched above by grasping both of its horns; we can preserve what is distinctive about an agent's self-knowledge without reducing its subject matter to a restricted range of inner happenings.

This paper explores how Anscombe's conception of non-observational knowledge can be put to work in explaining how the knowledge of one's own intentional actions differs both from the knowledge that others can have of those actions and from the knowledge a person has of what he *unintentionally* does. It also argues that several other recent attempts to work out this sort of position fall short of their goal by mischaracterizing certain key features of intentional actions and the kind of knowledge agents have of them. I begin in Section 2 by sharpening the central question, considering some recent attempts to address it, and arguing that none is fully successful. Section 3 then argues for an alternative position, according to which it is not because it has a special source but rather because of the distinctive role that it plays in our mental and behavioral economy that the knowledge we have of our own intentional actions is epistemically distinctive, and that such knowledge deserves to be called non-observational even though sense perception often plays a fundamental role in making one aware of what one does.

2 The Self-Knowledge Problem

How should we think about the way in which a person's knowledge of his own intentional actions differs from whatever knowledge he might have of other events, especially his unintentional behaviors and the intentional and unintentional actions of others? One familiar answer is that this holds because only knowledge of the former sort is essentially independent of *observation*, and a standard assumption in the literature on this idea is that the notion of non-observational knowledge is best understood as picking out a kind of knowledge that is somehow obtained independently of sense perception. Thus Kieran Setiya writes that "when an agent is

doing φ intentionally, he knows that he is φ -ing, and he knows this spontaneously, not on the basis of empirical evidence” (2008: 392); and in general, though this assumption is rarely made as explicit as Setiya makes it here, the idea that independence from observation equals independence from the deliverances of sense perception is shared by defenders and critics of this position alike.

The *locus classicus* for the idea of non-observational knowledge is the work of G.E.M. Anscombe, who frequently tries to bring out the distinctiveness of agents’ knowledge of their own intentional actions with examples like the following:

Say I go over to the window and open it. Someone who hears me moving calls out: What are you doing making that noise? I reply ‘Opening the window’. I have called such a statement knowledge all along; and precisely because in such a case what I say is true—I do open the window; and that means that the window is getting opened by the movements of the body out of whose mouth those words come. But I don’t say the words like this: ‘Let me see, what is this body bringing about? Ah yes! the opening of the window’. Or even like this: ‘Let me see, what are my movements bringing about? The opening of the window’. To see this, if it is not already plain, contrast this case with the following one: I open the window and it focuses a spot of light on the wall. Someone who cannot see me but can see the wall says ‘What are you doing making that light come on the wall?’ and I say ‘Ah yes, it’s opening the window that does it’, or ‘That always happens when one opens the window at midday if the sun is shining.’ (2000: 51)

As we noted above, the standard way of interpreting an example like this one is to say that in the second case it is because the agent has to *perceive* that she is casting the light on the wall in order to know that she is doing it that casting the light cannot be counted among her intentional actions, while in the first case her action is intentional because she doesn’t rely on sense perception in this way to know that she is opening the window. It is for this reason that,

according to Anscombe, replying “I knew I was doing that, but only because I observed it” to a request for a reason-giving account of an action (as for example “if one noticed that one operated the traffic lights in crossing a road”), is a way of implying that that action was not intended (2000: 14). Having to “observe” or “notice” that you are doing something in order to know that you are doing it rules out the possibility that it is among the things you intentionally do.

Nevertheless Anscombe is acutely aware of the intuitive difficulty in understanding how things like my opening a window can be among the things I know in such a special way, and she anticipates the following objection:

‘Known without observation’ may very well be a justifiable formula for knowledge of the position and movements of one’s limbs, but you have spoken of all intentional action as falling under this concept. Now it may be e.g. that one paints a wall yellow, meaning to do so. But is it reasonable to say that one ‘knows without observation’ that one is painting a wall yellow? And similarly for all sorts of actions: any actions that is, that are described under any aspect beyond that of bodily movement. (2000: 50)

We will consider soon enough how this problem might be addressed, but for now we should note that this case is by no means singular, and there are lots of different behaviors that suggest concerns along these lines: painting a wall yellow is one; opening a window and casting a light on a wall are two more; and so are reading a magazine, lifting a box, chopping an onion, putting ham on a sandwich, and so on. In each case the prima facie difficulty with treating an agent’s knowledge of such actions as non-observational arises from the fact that (1) it seems impossible actually to engage (and not merely: try to engage) in an action of the type in question without effecting changes in the world that involve more than the movements of one’s body, but (2) it appears that an agent will often know only by sense perception that he is bringing such extra-bodily changes about, and so (3) as a consequence, it seems that in cases like these an agent will

often know only by sense perception what he is intentionally doing. Anyone who wants to use independence from sense perception to explain the privileged access we seem to have to our intentional actions therefore needs to reject one or both of premises (1) and (2), or explain how the inference from these premises to the conclusion in (3) is invalid.

Call the problem of explaining how, in cases like these, agents can know without observation whatever they are intentionally doing the *Self-Knowledge Problem* for the theory of intentional action. Later on in this paper I will articulate a conception of non-observational knowledge that avoids falling prey to the Self-Knowledge Problem, but first I want to consider a few other responses to it and show why each of them fails.

2.1 *Rejecting (1)?*

According to the first of the premises that led us to the Self-Knowledge Problem, there are some actions it is impossible to engage in without bringing about changes in the world at a distance from one's body; and consequently it is impossible to know whether one is acting in such a way without knowing whether such extra-bodily happenings really are taking place. It is this idea, which Anscombe illustrates with the slogan "I do what happens", that can lead to the second horn of the dilemma sketched in Section 1, according to which human action is usually more than a mere interior turning of the wheel.

Yet there are some cases in which our commonsense metaphysics of action allows us to think of people as engaged in certain actions even when they are *not* bringing about the sorts of worldly changes that would be required for their completion: for example, I can spend many years working on a monograph even as nothing that deserves that title manages to materialize; and moreover I can be said to be working on it at those times when I am not actively engaged in

any real “work” at all, even when the status of my writing project could not be further from my mind. Thus Anscombe writes:

A man can *be doing* something which he nevertheless does not *do*, if it is some process or enterprise which takes time to complete and of which therefore, if it is not cut short at any time, we may say that he *was doing* it, but *did not do* it. This point, however, is in no way peculiar to intentional action; for we can say that something was falling over but did not fall (because something stopped it). Therefore we do not appeal to the presence of intention to justify the statement ‘He is Y-ing’; though in some cases his own statement that he is Y-ing may, at a certain stage of the proceedings, be needed for anybody else to be able to say he is Y-ing, since not enough has gone on for that to be evident; as when we see a man doing things with an array of wires and plugs and so on. (2000: 39)

Noting this feature of our action concepts can incline us to think that whether or not one acts in a given way doesn’t actually have to hinge on the kinds of worldly changes that premise (1) demands: if I can act in a given way in the absence of the kinds of extra-bodily events that would constitute a completed act of the sort in question, then it seems that I don’t need to have knowledge of such events in order to know what I do.¹ This claim is compatible with the idea that one very often has to accomplish *something* extra-mental in order to act, but perhaps the idea is it is possible to engage in a certain kind of action, say one of ϕ -ing, not just by bringing about the happenings that a *completed* act of ϕ -ing involves but also by doing something else, say ψ -ing, that is a *means* to such a completed act, where one’s ψ -ing can consist of happenings that take place at less of a distance from one’s body than would those that constitute the fully completed act of ϕ -ing.² And since (one might suppose) it is possible to have non-observational knowledge that one is acting in the “proximal” manner, we can know in this way that we are doing all of the various things we are intentionally up to.

But such a position runs dangerously close to what Anscombe calls the “false avenue of escape” according to which “I really ‘do’ in the intentional sense whatever I think I am doing” (2000: 52). According to Anscombe the problem with this way of thinking is that it is impossible to say how the “vehicle” for an intention could be independent in this way of a worldly occurrence that is the achieving of what one intends; but whether or not we are moved by such a response, this objection to premise (1) is prey to a more immediate problem, namely that ordinary intuitions about what it takes to engage in an action are not nearly as liberal as the objection supposes. As Kevin Falvey notes:

The precise details of the truth-conditions of sentences in the progressive [tense] is a matter of considerable obscurity, but it is surely a necessary condition for me to be φ -ing that it be possible for me to φ . Hence the claim that I am or was φ -ing must be withdrawn if in the course of things it becomes clear that I cannot φ ; either because I lack the general skills or ability to φ , or because I lack the materials required for φ -ing. If I put the kettle on the stove and say, “I’m making tea,” and am told that the stove isn’t working, I should not subsequently say that I was *making tea* when I learned that the stove wasn’t working, I would retreat to something like, “I was going to make tea” (or, “I was trying to make tea”). On the other hand, if, as I am putting the kettle on for tea, the phone rings and the tea-making is aborted, I could still later say that I was making tea when my friend called, as a result of which I didn’t make tea. The openness of the progressive allows for interruptions of actions-in-progress, including changes of mind. But from the fact that an event or process of a given type could not have been completed in the circumstances, it does seem to follow that no event of that type could have been underway. (2000: 24)

Falvey is right, and in fact the point applies more widely than those instances where the sorts of actions under consideration are simply impossible: to return to our original example, it seems that a person who means to paint a wall yellow but is accidentally using white paint instead simply doesn’t count as painting the wall yellow on our ordinary understanding of what such an

activity involves; he's *trying* to paint the wall yellow, we might say, and while he's doing some of the things that are required actually to do that it's nevertheless not the action he's really engaged in. So premise (1) seems well enough in order: there are plenty of cases where acting intentionally in a certain way *does* require bringing about specific changes at a distance from one's own body, and the crucial question concerns whether an agent will always know without observation that such changes are being brought about.

2.2 *Rejecting (2)?*

If human agents essentially possess non-observational knowledge of their intentional actions even under descriptions involving changes that take place at a distance from their own bodies, and if what it is for something to be non-observationally known is for it to be known independently of sense perception, then it seems that whenever someone is intentionally engaged in an action that requires bringing about certain changes in the world he must know independently of sense perception that he is bringing such changes about. Attempting to explain how such knowledge is possible is the most common way that defenders of a broadly Anscombean theory of action have responded to the Self-Knowledge Problem, and this section considers a couple such attempts and argues that they fail to show that our premise (2) can plausibly be rejected.³

2.2.1 *Knowledge by "Know-How"*

In his paper "Practical Knowledge", Kieran Setiya that the Self-Knowledge Problem vanishes once we recognize two things: first, that the intention to act in a certain way consists partly in the *belief* that one will do just that; and second, that when we intend to do things we generally *know how* to do them, and moreover know that we have such know-how. According to

Setiya, once these factors are in place there is no trouble in seeing how an exercise of human agency essentially involves non-observational knowledge of what one intentionally does: for assuming that an action is intentional if and only if the agent intends to engage in it, then the first principle yields the claim that the agent will also *believe* that he is engaging in it; while the second yields the further conclusion that this belief will be justified by his knowledge that he'll be able to do it if he so intends. As Setiya puts it:

Knowing how to φ is the state or condition that, with knowledge of ability, provides the epistemic warrant for decision. Together, they justify the transition in which one forms the intention and belief that one is doing φ or that one is going to do it. More carefully, this transition is justified if and only if one's decision is an exercise of knowledge how to φ and one has knowledge of ability, in the simple conditional sense. Knowledge how thus plays a role in dynamic epistemology, in our entitlement to form and revise beliefs. (2008: 407)

By “the simple conditional sense” in which we have knowledge of our abilities Setiya means the sense in which a person can know that, if he intends to be doing something at some time, then at that time he will be doing so in fact. If we possess knowledge of this sort whenever we act intentionally, and if such an action will always be accompanied by a justified belief in its occurrence, then will it follow that our knowledge of what we intentionally do is essentially independent of sense perception?

But the fundamental difficulty with this position is simply that it leaves untouched the possibility that even a perfectly able agent will sometimes *fail* to do what he intends, and when the possibility of such failure is sufficiently real we are left without a way to understand how the agent knows that such a possibility does not obtain. Even Anscombe, despite insisting that “the failure to execute intentions is necessarily the rare exception” (2000: 87), allows that intentions *sometimes* “fail to get executed” (2000: 82); and clearly it is possible to act intentionally in a

given manner even in instances where, one's general ability so to act notwithstanding, the surrounding circumstances raise a real possibility that one's intended action might fail to materialize. If, for example, I am presently painting a wall a perfectly uniform shade of yellow, then even though I may (i) believe that I am painting it in such a way and (ii) know that I am *able* so to paint it, nevertheless I may not (iii) *know that that is the way I am painting it* without seeing how the paint is going on – for after all, getting the paint just right can be a pretty tricky business. Having the ability to act in a certain way cannot require being able to do so successfully *whenever* one sees fit to try; and even if we did read Setiya's account in this way then it would fail to apply to the vast majority of our intentional actions, since of course we're not generally able to realize our intentions with a perfect a rate of success. It may be that intentional agency always involves a non-sensory knowledge of what one is *able* to do, but that alone is not sufficient always to justify us in believing that we have followed through on our intentions, and so Setiya's account does not avoid the Self-Knowledge Problem after all.

2.2.2 *Inference from Intention*

In "How We Know What We Are Doing", Sarah Paul argues that we have knowledge of our intentional actions not by the evidence of sense perception but rather thanks to an (often unconscious) *inference* from the knowledge of what we intend, plus empirical and non-empirical knowledge of the background conditions that make success likely. Thus "what the agent knows evidentially", Paul writes, "is what he intends to be doing, while insofar as he has a belief about what he is actually doing, this is evidentially based on his knowledge of what he intends, plus his evidence for thinking that he will do what he intends" (2009: 12). Thus unlike Setiya, Paul rejects the idea that intending to act in a certain way *essentially* involves the belief that that is how one will act; she shares his commitment, however, to the thought that the justification for an

agent's belief that he is acting intentionally in a given way is essentially independent of any perceptual awareness of that action.

But Paul is sensitive to the sort of objection raised just above against Setiya's view, that merely knowing that one *can*, even most of the time, act in a certain way if one intends to will not always justify one in believing that he or she *is* so acting on any given occasion. Thus she allows that the agent's "background knowledge of his circumstances as being conducive to his ϕ -ing – or at least, [of] the absence of reasons to believe his ϕ -ing will be obstructed" will be a part of what grounds his knowledge of what he does (2009: 15): *such* awareness may be observational, of course, but Paul claims that since it "is not experience or observation *of the particular action in question*" (2009: 16; emphasis added), it will not be the sort of perception-based knowledge that Anscombe's account is meant to rule out. According to Paul, it is because the knowledge of one's intentions and abilities is accompanied by the knowledge of one's circumstances that one is entitled to infer what one is doing, and that the conclusion of such an inference will always count as knowledge: given that I know (i) what I presently intend to do, (ii) that I usually do what I presently intend, and (iii) that I am capable of doing just that in circumstances that (iv) I know to obtain, I am justified in believing that I am presently acting just as I intend to be.

Is Paul's inferential view a viable response to the Self-Knowledge Problem? One question is whether Paul fully preserves the first-personal/third-personal asymmetry that was emphasized at the beginning of this paper, i.e. the way in which an agent's knowledge of his or her own intentional actions seems essentially different from any outside observers' knowledge of them. For on an account like Paul's while there will clearly be fundamental differences between your and my respective ways of knowing about my *intentions*, it appears that once that

knowledge is in place you may be in every bit as good a position as I am to know what I am doing on the basis of the knowledge of those intentions plus my general tendencies and the evident favorability of my circumstances: this is an inference that you are every bit as capable of making as I am, and it seems that you will be no less justified in it. But it seems as if it is knowledge of one's *actions*, and not just of the intentions that underlie them, that ought to be characterized by first-personal privilege, and it is a significant defect of Paul's position if it fails to ensure this.

Moreover, on careful inspection Paul's account does not fare much better than Setiya's when we consider the question how we know what we are doing in sufficiently *unfavorable* circumstances. If, for example, I am successfully steering my boat through a narrow passage, my perceptual awareness of the storminess of the sea may give me reason to believe that the surrounding conditions are not especially conducive to steering the boat as I intend to. In such a case, to know that I am steering in the intended direction rather than (unintentionally) heading for a dangerous pile of rocks I will need, therefore, to rely on my perceptual experience of *where I am going*: yet surely my going the direction I am can nevertheless be an intentional action. If relying on sense perception in this way to know that I am acting (and not merely trying to act) as I intend entails that my knowledge of what I am doing is observational, then the claim that we *always* know our intentional actions without observation is clearly false. It may be that Paul's account is sufficient to explain how we can be justified in believing that we are acting as we intend some or even much of the time, but it cannot apply as generally as we should want it to. The Self-Knowledge Problem remains.

2.3 *Rejecting the Inference to (3)?*

Once we accept premises (1) and (2) in our argument for the Self-Knowledge Problem it might seem that the conclusion in (3) follows immediately; yet Richard Moran has staked out a position that can be read as an attempt to deny just this. According to Moran, what allows an agent's knowledge of his own intentional actions to "involve reference to actual changes in the world" beyond the confines of his body while still differing from an outside observer's "speculative, observational knowledge of what is the case", is that the agent's knowledge "commits itself not only to the obtaining of certain events in the world, but to the specification of the descriptions under which what happens counts as the execution of his intention" (2004: 56). The guiding idea here is that our understanding of intentional action is essentially an *intensional* affair: "to say", as Anscombe puts it, "that a man knows he is doing X is to give a description of what he is doing *under which* he knows it" (2000: 12). And Moran's proposal is that the difference between the non-observational knowledge of one's own actions and the observational knowledge of unintentional behaviors or the actions of other agents is that an agent's knowledge of the proper intentional *characterizations* of his actions is essentially of the former sort:

Understood extensionally, I can know *what happens* only through observation, including the perception that serves as an aid in the execution of action such as writing on the blackboard ..., and the causal knowledge Anscombe refers to earlier as "knowledge or opinion concerning what is the case, and what can happen—say, Z—if one does certain things, say ABC" ... With this empirical knowledge in place, I can form an intention to *do* something, such as opening a window, and then actually *do that thing*. But the event which is my action only counts as something I intentionally do in virtue of some of its descriptions and not others, and my knowledge of it can be said to be 'non-observational' *only* under the terms of such descriptions. (2004: 55-56)

So unlike Paul and Setiya, Moran seems to allow that my knowledge that I am actually effecting the sorts of changes in the world that are required for me to count as acting as I intend will sometimes involve a sense-perceptual awareness of my bringing certain things about. Yet he wants to resist the idea that this means my knowledge *of my actions* can't meaningfully be classified as non-observational, since my "empirical knowledge" of what happens in the world won't be enough to settle the descriptions under which those happenings count as intentional behaviors. This latter sort of knowledge "depends on, but does not reduce to, the speculative knowledge of what can happen and what is happening, and in this way Anscombe may evade the charge ... that the admitted dependence of successful action on ordinary observation must mean that the agent's knowledge of what he is doing is ultimately observational after all" (Moran 2004: 56).

But is it really true that "observational knowledge of what is the case" essentially fails to specify the description or descriptions under which an observed happening counts as an intentional action? Moran motivates this idea by appealing to an example of David Velleman's (1989: 15), in which a person walking through Manhattan suddenly realizes that he doesn't know what he is doing. Of this example, Moran writes:

... such cases enable us to see how it can be true both that the agent will normally depend on observation of various kinds to carry through an action like walking up Fifth Avenue ..., and also that the knowledge that is temporarily lost in such cases is nonetheless not made up for my further observations alone. The person stops and *looks around*, observing his position and his environment for clues as to what he might have been up to, but this by itself does not deliver to him the knowledge of what he is or was doing, for it does not provide him with the particular set of descriptions of what he sees or the movements he is making, under which what he was doing counted as an intentional action of his. (2004: 57)

This example is supposed to help us see, Moran continues, “how practical knowledge *could not* be observational, could not be perceptually derived from the world. For nothing the agent sees in the world could give him those descriptions, even though *what* is claimed in practical knowledge is a world-involving matter of fact” (ibid.). Yet this conclusion – that the knowledge of what one is intentionally doing requires a component which observation cannot possibly supply – only follows from Velleman’s example if the content of the walker’s perceptual experience *is the same* in both these cases, i.e. when he is walking absent-mindedly and without knowledge of what he is doing, and when he is walking intentionally and with full self-knowledge: and it is only if we make the implausible assumption that an agent perceptually experiences his intentional and absent-minded behaviors alike only as sets of bare bodily movements that we would be entitled to treat these two cases in such a way.⁴ Moran is certainly right to say that “further observation of the world” (2004: 58) will not provide the man who stops and asks what he is doing with the right sort of answer to the question of what he is doing, but it does not follow from this alone that “the knowledge he has temporarily lost is something that goes beyond his observational knowledge” (ibid.), i.e. that when he *did* know what he was doing, this knowledge had a non-empirical component. This would be a valid inference only if what was in the world to be observed was the same when he was acting intentionally as when he was not, but clearly there is no good reason to think that this is so. Indeed, in the very same paper, Moran himself seems to suggest that it is not:

There is an ordinary sense in which [the absent-minded walker] *sees* the same things as he did when he was engaged in action (e.g., walking up Fifth Avenue). But there is also a sense in which what he sees is now blank to him, because he cannot articulate his own relationship to it. The objects and scenes of his environment no longer have a role assigned to them in his ongoing action (as goal, obstacle, distraction, background, etc.). He sees Fifth Avenue, and he sees it from a

particular perspective which indicates what direction he was heading in, but he does not see his goal. Instead, he sees *everything* around him, the traffic going by, the meaningless street signs, the strangers' faces, and nothing in those details enables him to discern a destination, a point to his being right here facing in this direction. (2004: 57)

If what Moran says here is right, we are perfectly entitled to say that the reason sense experience does not provide the absentminded walker with knowledge of his intentional actions is that in his case there are no such actions to be experienced at all, while for the person presently exercising intentional agency, things are experientially quite different, as his intentional behaviors are right out in the open.

Not only is Moran's account insufficiently motivated, but it is also one on which an agent's self-knowledge is non-observational only in a rather weak sense, one which runs perilously close to a "two-factor" view according to which the only thing *really* known by the agent without observation is something like what he intends.⁵ For if an agent knows without observation only the *descriptions* under which his actions counts as intentional, whereas he needs to rely on observation to become aware of the worldly happenings that fall under such descriptions, then it is hard to see how it is the action *itself*, and not just some aspect of it, that is non-observationally known. By contrast, what Anscombe means to point us toward is a view according to which *even what happens* in a case of intentional action is something that an agent knows about in a special sort of way: it is not just that I will always know in a way that you cannot the descriptions under which my bodily behaviors count as intentional actions, but also that I will have a special epistemic relationship to those happenings even in the "extensional" sense, even as they involve "mere" changes in the physical world. It is this that Anscombe seems to be after when she writes that "in so far as one is observing, inferring etc. that [some event] Z is actually taking place, one's knowledge [of Z] is not the knowledge that a man has of his

intentional actions” (2000: 50); and it is hard to see how this holds on Moran’s analysis. In the next section, I will argue that it is by developing a conception of non-observational knowledge that does not require independence from sense experience that we can avoid the Self-Knowledge Problem, and account for the essential epistemic distinctiveness of the knowledge of one’s intentional actions without taking on any unacceptable commitments.

3 Practical Knowledge

According to the argument of Section 2 it is impossible to reconcile our commonsense metaphysics of human action with the idea that our knowledge of what we intentionally do is essentially independent of sense perception: there are, we saw, plenty of possible cases where an agent’s perceptual awareness of his actions may play an important role in grounding his knowledge of what he is doing, without making the action in question *ipso facto* unintentional. But how can this conclusion be squared with the deeply intuitive idea that an agent’s knowledge of his intentional actions is epistemically distinctive, differing both from his knowledge of other sorts of things and from whatever knowledge of his actions might be had by an outside observer?

Well, what exactly do we mean in this context when we speak of an *observer*, and of knowledge had *by observation*? Surely the implication of a reliance on sense perception is important here, but the notion of observation connotes more than this: observation involves not just perceiving something, but doing so somehow *passively*; to observe something is to sit back, as it were, and simply take it in for whatever it happens to be. Brian O’Shaughnessy highlights this aspect of our concept of observation when he argues for the incoherence of self-observation in *The Will*:

The conductor cannot listen as *observer* to the music he makes, since he already listens to it from the standpoint of *creator*. His listening is logically subordinated to that activity, to which it relates somewhat as the painter's looking relates to the activity of painting. Were he to listen as observer, his listening would no longer be co-ordinated with and logically subordinated to the act of making music. In sum: one cannot be listening to the music one makes both from within the act and from without the act; one cannot simultaneously listen in two different ways. Now this is the nature of the difficulty where the putative observer-sense is playing an essential stage-setter role for the activity it putatively studies. (1980: 29)

Clearly O'Shaughnessy is not denying the obvious fact that in a case like the one he describes the conductor can and will *hear* the music his orchestra makes; indeed, in most cases it will be only because he does hear it that he can go on conducting, which is what O'Shaughnessy means by speaking of the "essential stage-setter role" that the conductor's auditory perception plays in his ongoing activity. Rather, his claim is that there are certain *ways* of listening or otherwise attentively perceiving that are impossible when the object of perception is the perceiver's own action: in taking such a dual attitude toward what he does, O'Shaughnessy says, the (putative) agent-observer would have to place himself both "within" and "without" the very same action (1980: 31-32), and this is a stance no person can possibly adopt.

This way of marking the distinction between the kinds of action awareness characteristic of agency and observation respectively is illustrated very schematically in (Figure 1). On my reading, O'Shaughnessy's point in denying that a person can ever observe his intentional actions is not that human agents are never sense-perceptually *aware of* the things they intentionally do (this is the point of the downward arrow from action to awareness on the "Agent" side), but rather that the kind of self-awareness characteristic of intentional agency is essentially bound up in the activity of self-*control*. What makes the observational mode fundamentally passive is that,

while it may exist alongside a preference for things to turn out one way rather than another (think of a fan’s perspective on a sporting event), the observer is not the one who brings those changes about; and thus the experienced world of the observer “is one that is going its own way, is taking its own course” (O’Shaughnessy 1980: 20). By contrast, an agent keeps track of his actions not only with a purpose in mind but also in such a way that he takes it to be at least partly up to him to have things turn out accordingly; the agent is a “creator”, and his awareness of his actions is always subordinated to the task of ensuring that things proceed along the course he intends.

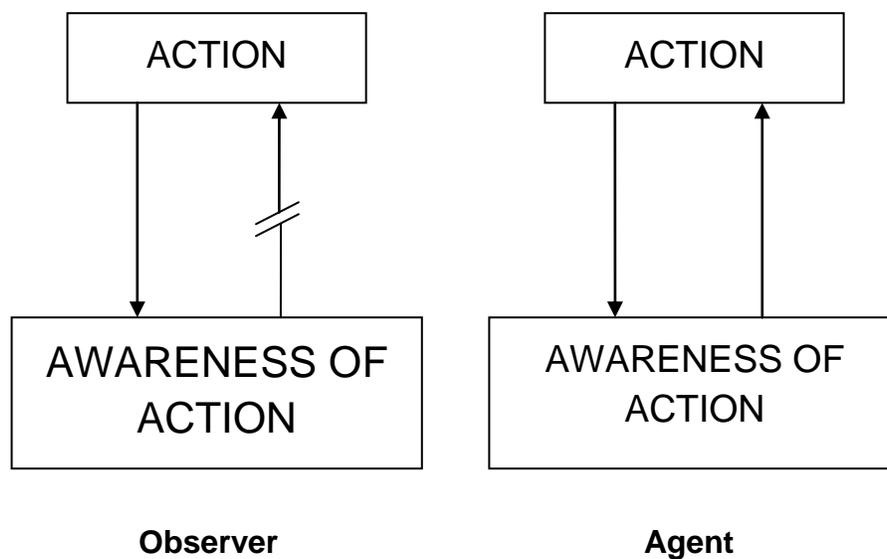


Figure 1. O’Shaughnessy’s distinction between the agent and observer models of action awareness.

For our immediate purposes, the crucial point to note right now is that if the distinction between observational and non-observational knowledge breaks down along these lines rather than hinging just on whether the knowledge in question is perceptually grounded, then there is reason to hope that the Self-Knowledge Problem as articulated in Section 2 will be no problem for us at all: if knowledge can be non-observational in virtue of something other than its source, we will be able to reject the assumption that just because we sometimes have to rely on our

sensory awareness of our actions to know that we are doing as we intend, the knowledge of our actions therefore requires us to go in for self-*observation*. The remainder of this section will develop and defend this conception in a bit more detail, suggesting some important parallels between O’Shaughnessy’s discussion of self-observation and Anscombe’s notion of “practical knowledge”. On this account, it is ultimately because of its functional role, rather than its source, that an agent’s knowledge of his intentional actions counts as non-observational, and is characterized by the epistemic asymmetries that we have been trying to account for.

3.1 “The Cause of What it Understands”

O’Shaughnessy’s discussion of self-observation locates the distinguishing feature of an agent’s self-awareness not in its sources, but rather in the distinctive causal relationship that it has to its subject-matter. To emphasize this aspect of the awareness of one’s intentional actions is not, of course, to deny that that self-awareness very often *has* distinctive sources: for a human agent is, at least ordinarily, aware of the movements of his body not just through ordinary sense perception but also through intra-bodily channels that give him a distinctive way of telling where his limbs are at any given moment.⁶ But Anscombe’s account of intentional agency invites us to reject the idea that the involvement of such privileged sensory channels in an agent’s awareness of his actions could constitute the *entirety* of what makes that awareness distinctive; and thus she gives us cases like the following:

... suppose someone simply wanted to produce the effect that in fact I lowered my arm at the speed at which it would fall—he is a physiologist, and wants to see if I generate anything different in my nerve fibres if I do this. So he fixes up a mechanism in which something in motion can be kept level if I hold a handle and execute a pumping movement with my arm and on the downward stroke lower it at the rate at which it would fall. No my instruction is: Keep it level, and with a bit

of practice I learn to do so. My account of what I am doing is that I am keeping the thing level; I don't consider the movement of my arm at all. I am able to give a much more exact account of what I am doing at a distance than of what my arm is doing. So my keeping the thing level is not at all something which I calculate as the effect of what I really and immediately am doing, and therefore directly know in my 'knowledge of my own action'. In general, as Aristotle says, one does not deliberate about an acquired skill; the description of what one is doing, which one completely understands, is at a distance from the details of one's movements, which one does not consider at all. (2000: 54)

Call this case the *Pump Example*. On Anscombe's analysis of the Pump Example, when she keeps the thing level she must have non-observational knowledge that this is what she does. Yet clearly her knowing this cannot be explained just in terms of her reliance on bodily proprioception in what Shaun Gallagher calls "the ordinary (non-visual) sense of somatic (mechanical) information about joint position and limb extension" (2005: 46): for while there are *some* aspects of what she is doing that Anscombe knows about in this sort of way, knowing that one is moving one's body is not itself a way of knowing that one is keeping something level, and it is only under the latter description that Anscombe's action is intentional. Given this, the Pump Example is one where the difficulty of the Self-Knowledge Problem becomes especially acute: the action of keeping the thing level is one that requires Anscombe to achieve a certain effect in the extra-bodily device, and if the task is sufficiently difficult then her knowledge of that effect will be based on her visual awareness of it. But given O'Shaughnessy's account of the observational/non-observational distinction this is no threat to the idea that Anscombe knows about her action without observing it: to be sure, she is *aware* of what she does, but what makes this awareness non-observational is not its source but rather the role it plays in shaping the unfolding of the very action it is an awareness of.⁷

It is not, however, only the agent's *awareness* of his action that has a special role to play in the way that action unfolds; for according to Anscombe there is also a way in which our *knowledge* of our own actions is causally relevant in bringing them about. In developing this conception of "practical knowledge" she references a passage from the *Summa Theologiae*:

... man is more like God with respect to his practical intellect, which is the cause of things thought of, than his speculative intellect, which derives knowledge from things. Therefore man's happiness consists in activity of his practical rather than his speculative intellect. (Aquinas 1983: 34)

Aquinas goes on to reject this account of human happiness as giving insufficient weight to the nobility of speculation, though he agrees that "the practical intellect is related to what it knows as God to what He knows" (1983: 35): that is, he holds that human agents are knowers not just by the ways we conform our minds to the world but also in the knowledgeable ways we bring things about. This is why in the Pump Example, when Anscombe keeps the device level by operating the handle we can explain her actions partly in terms of her knowledge of how those very actions unfold: it is *because* she knows that she is keeping the thing level that she proceeds along as she has been; whereas if she comes to understand that the thing is off kilter she will reshape her action accordingly.

Yet as our analysis of the Self-Knowledge Problem showed us, human agents are quite *unlike* God in that it is not simply by saying the word that we can do what we will; and therefore it is generally only in virtue of the appropriate sort of sensory awareness of his actions that an agent's "practical knowledge" of what he intentionally does will be fully in place. Not all of the sensory feedback relevant to behavioral control proceeds by way of what we *know*, of course: this is one of the lessons of the example cited just above, where the agent operating the pump

handle “completely understands” only whether the device is being kept level, and in doing this she makes tacit use of proprioceptive feedback but takes little or no explicit account of how she is moving her limbs. But as we noted, even in this case there is *something* which the agent knows about, namely the status of the thing she is trying to keep level, and it is partly because that empirical-but-thoroughly-non-observational knowledge is in place that her action unfolds in the particular way it does.

Thus as the picture of agent-awareness sketched in Figure 1 already indicated, on this account an agent’s practical knowledge is *itself* a kind of knowledge that is “derived from things”, even if it is not “speculative” in the ordinary sense of the term: we are like God in being able intentionally to bring things about, yet thanks to our finitude we generally have to rely on sense perception both in getting things done and in coming to know whether or not we are acting as we intend. As we saw, it is only very rarely that a human agent intentionally engages in an action so simple that he can successfully achieve, much less know that he achieves, what he intends without paying any perceptual attention at all to what he does. In less simple cases, the agent’s awareness, and thus also the agent’s knowledge, of his behaviors and their effects on the world is part of a feedback loop that shapes and is in turn shaped by the ways his actions unfold. It is because the conductor can *hear* the music that the instrumentalists produce that he knows how things presently stand and where he needs to proceed from here; and if his knowledge were other than it is – if, say, the musicians got off track, or if he misheard which notes the strings were playing and felt a mistaken need to correct them – then the future course of his action would be quite different. The agent does not act without keeping track of his actions; he cannot be knowledgeably self-controlling without being perceptually self-aware. In this way human agents are at once doers and perceivers, in a way that the divine Agent is not.

This account of the non-observational, “practical” knowledge of one’s own intentional actions is illustrated very schematically in (Figure 2). What it proposes is that the epistemic privilege characteristic of the knowledge of one’s own intentional actions should be explained not by the fact that that knowledge has a special source, but rather by the special functional relationship that that knowledge bears to the action that is known: We might put the point by saying that it is because an agent’s knowledge of what he is intentionally doing is essentially a *part* of the action itself⁸ that he knows about that action in a way that an outside observer cannot. This view preserves the first-person/third-person asymmetry that is fundamental to the knowledge of our intentional actions without denying that sense perception is implicated in the way we know what we are doing.

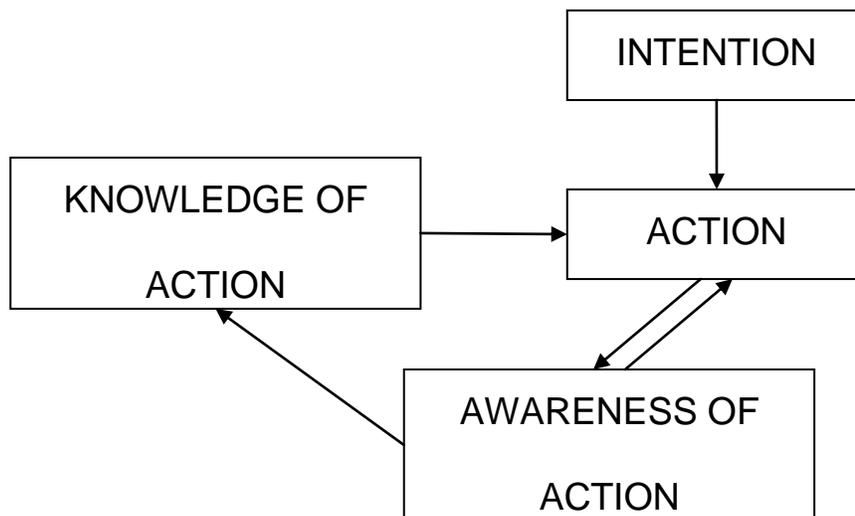


Figure 2. Non-observational, practical knowledge of one’s intentional actions.

3.2 Intended and Unintended Actions

When we discussed the epistemic distinctiveness of an agent’s self-knowledge in Section 1, we suggested that one’s knowledge of one’s own intentional actions differs fundamentally both from an outsider’s knowledge of them and from whatever knowledge one may have of one’s own

unintended behaviors. The argument of Section 3.1 focused on the first of these asymmetries, but what of the latter? To illustrate this distinction Anscombe describes a case in which “one noticed that one operated the traffic lights in crossing a road”, such that saying “I knew I was doing that, but only because I observed it” would be a way of marking the operation of the lights as unintentional (2000: 14). How should the notion of practical knowledge that I have been developing here lead us to think about this second asymmetry?

The natural answer is that it is because of the integral relationship between intentional action and the agent’s knowledge of it that it does not make sense for an agent to have to “stop and look” to determine what he is intentionally doing: if this sort of thing were to take place, then the action in question could not have been one of which the agent had a genuinely practical knowledge. On this account the man in Anscombe’s example has practical, non-observational knowledge of his crossing the street but merely speculative, observational knowledge of his operating of the lights not because he knows about the first activity independently of perception, but rather because only this first piece of knowledge is “the cause of what it understands”, in the sense that that knowledge plays a role in keeping his behavior on course. If the man comes to know (or think) that he isn’t crossing the street after all but rather veering off into traffic, say, then he will try to adjust his behavior so as to do the thing he intends; while so long as he knows that he’s crossing the street, then all else being equal⁹ he will proceed contentedly along. By contrast the man’s knowledge that he is operating the traffic lights has no such status: it is knowledge that is entirely extrinsic to the behaviors in question, as the man has no stake in whether he operates the lights or not, and indeed if it seemed to us as if he *did* have such a stake in ensuring that he was operating the lights then we’d also find ourselves believing that that was not an unintended behavior at all. It is, as Velleman’s discussion of the example of the absent-

minded walker suggests, possible to be aware of certain of one's own actions from the perspective of passive observer rather than deliberately self-controlling agent, and so to have a knowledge of them that is speculative rather than practical; it is just that when this is the case, the actions in question are *ipso facto* unintentional ones.

Thus once again, the crucial contrast underlying this asymmetry is not between perception-based and perception-independent ways of knowing about an action but between knowledge that is causally implicated in shaping and sustaining the very action it is knowledge of, and knowledge that lacks this distinctive causal role. What the man crossing the street indicates when he says that he knows himself to be operating the lights “only because [he] observed it” is not simply that he relied on sense perception in coming to know this about himself, but also that the knowledge in question is ultimately “accidental” to the action in the way that the knowledge of an intentional action never can be. From the perspective of the agent, that he has been operating the lights (or: casting sunlight on the wall, etc.) is something that he just discovers about himself in the same way that *we* might have happened upon his doing that, while his knowledge that he was walking was quite different from this. With respect to his action under the former description, but not the latter, he simply “caught himself in the act”, and – unless perhaps he *then* decided that operating the traffic lights was something of a lark and so began to do it intentionally at that point – his knowledge of whether he was engaged in that action had nothing to do with whether the relevant events took place.¹⁰ This is just what it is to be an observer rather than an intentional agent, a passive taker-in of facts rather than a deliberate bringer-about.

Acknowledgements

I am grateful especially to Randall Amano, John Campbell, Ben Kiesewetter, Niko Kolodny, Thane Naberhaus, Alva Noë, and several anonymous referees for helpful feedback regarding various aspects of this material.

Notes on Contributors

John Schwenkler received his doctorate from U.C. Berkeley in December, 2009. Since then he has been Assistant Professor of Philosophy at Mount St. Mary's University in Emmitsburg, Maryland.

References

Anscombe, G. E. M. 2000. *Intention*, 2nd ed. Cambridge: Harvard University Press.

Aquinas, T. 1983. *Treatise on Happiness*. Trans. J. A. Oesterle. Notre Dame, IN: University of Notre Dame Press.

Davidson, D. 1980. *Essays on Actions and Events*. New York: Oxford University Press.

Donnellan, K. 1963. Knowing what I am doing. *Journal of Philosophy* 60: 401-09.

Falvey, K. 2000. Knowledge in intention. *Philosophical Studies* 99: 21-44.

Gallagher, S. 2005. *How the Body Shapes the Mind*. New York: Oxford University Press.

Gibbons, J. Seeing what you're doing. Forthcoming in *Oxford Studies in Epistemology*.

Hurley, S.L. 1998. *Consciousness in Action*. Cambridge: Harvard University Press.

Moran, R. 2000. Anscombe on 'practical knowledge'. In *Agency and Action*, ed. J. Hyman and H. Steward, 43-68. New York: Cambridge University Press.

O'Brien, L. 2003. On knowing one's own actions. In *Agency and Self-Awareness*, ed. J. Roessler and N. Eilan, 358-382. New York: Oxford University Press.

O'Shaughnessy, B. 1980. *The Will: A Dual Aspect Theory*, vol. 2. New York: Cambridge University Press.

Paul, S. K. 2009. How we know what we're doing. *Philosophers' Imprint* 9. At www.philosophersimprint.org/009011.

Proust, J. 2003. Perceiving intentions. In *Agency and Self-Awareness*, ed. J. Roessler and N. Eilan, 296-320. New York: Oxford University Press.

Roessler, J. 2003. Intentional action and self-awareness. In *Agency and Self-Awareness*, ed. J. Roessler and N. Eilan, 382-405. New York: Oxford University Press.

Setiya, K. 2008. Practical knowledge. *Ethics* 118: 388-409.

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

Notes

¹ For an example of such a response to the Self-Knowledge Problem, see Paul 2009: 16-17.

² There is a version of this latter idea in O'Brien 2003, who proposes that it is an agent's "inner" awareness of the bodily movements that constitute "basic actions" that non-observational knowledge consists. But Anscombe's example of the person operating a pump handle, which is discussed in Section 3.1 below, shows why this account cannot work: for in this case the agent is distinctly *unaware* of the details of her bodily movements; all she knows is whether she is keeping level a device whose position she visually perceives.

³ For a further objection to accounts like those that I discuss in the following sections, see Gibbons (forthcoming).

⁴ For criticism of such an assumption, see Proust 2003.

⁵ For one explicit statement of such a view see Donnellan 1963.

⁶ In forthcoming work, I criticize the idea that such bodily self-awareness should be understood as non-perceptual.

⁷ For some similar suggestions see Roessler 2003, which distinguishes the agent's use of perceptual attention to answer "practical question" from the observer's use of it to answer "theoretical" ones. Yet on Roessler's account the answer to a practical question is *not at all* a theoretical matter; its content is not a way of staking a claim as to how things are in the world. A more natural view, I think, is that the agent relies on perceptual experience to take in how things are and immediately implicates this knowledge in deliberately shaping those things in the intended ways; there is no reason why the relevant deliverances of sense perception cannot have a mind-world direction of fit. But further discussion of this point would take us too far afield.

⁸ This language is meant to recall the "two-way interdependence" account of action and perception developed in Hurley 1998.

⁹ "All else being equal" because it is possible to have multiple intentions, some of which can come into conflict with one another. For example, perhaps the man means to be crossing the street but also means *not* to be operating the lights, in which case if he knows that he's succeeding in one of these intentions but not the other he will have to make a choice as to which is more important to him.

¹⁰ We can say something similar about the question whether it is possible to act intentionally in a given way *without* knowing that one so acts – a possibility that Anscombe explicitly rejects (e.g. at 2000: 88, where she writes that in such a case "what happens does not come under the description—execution of intentions—whose characteristics we have been investigating"), but which has been brought out in a number of examples, e.g. Davidson's case of a clerk who makes a stack of legible carbon copies just as he intends even while he doubts whether he is actually pressing down hard enough (Davidson 1980: 50, 92). On this account, the thing to say about such cases is that an action can be intentional insofar as the agent is *poised* to bring any awareness he may gain of what he is doing into the ordinary cycle of intentional self-control. Even in a condition of relative ignorance, it may remain that *if* an agent had known that he *hadn't* been doing the things he intended to do, then he'd have been inclined to do something different to try to change that: thus Davidson's clerk is such that if he were made to *see* that he simply wasn't managing to produce all of the carbon copies he was after, then he'd begin pressing a bit

harder or at least expressing some dissatisfaction with his failure to get the job done; and if we came to believe that although he was in fact producing the whole stack of carbon copies he didn't possess this sort of disposition toward succeeding then we'd no longer think of him as making all of the copies intentionally after all. Gibbons (forthcoming) has a helpful discussion of what we should think about the significance of such cases for the possibility of privileged access to one's own actions.