8
Knowledgeable Action

8.1 Introduction

"Trust me, I know what I'm doing" was the catch phrase of Sledge Hammer, the star of the satirical 1980s cop show *Sledge Hammer!* The joke was that he didn't. *Urban Dictionary* makes the same joke. Its top entry for "trust me, I know what I'm doing" is this: "A phrase often used by people who don't really know what they're doing." This is kind of funny, because it's kind of true.

Now, human agents often know what they are up to. When they do, things are going pretty well for them. In this chapter I want to make sense of how knowledge contributes to a particular mode of agentive excellence—what I will call knowledgeable action.

Knowledgeable action is intentional action that embeds knowledge in specific ways. I say more below, but to preview: sometimes agents act intentionally. Sometimes when they do, they also act knowledgeably. Knowledgeable action is a mode of agentive excellence in extending the agent's success in ways I elucidate.

Two distinct targets of knowledge are of interest—knowledge of action and knowledge of how to act. Knowledge of action signifies the agent's knowledge that she is A-ing, as and when she is A-ing. Knowledge of how to act signifies knowledge states that are about ways to A.

My guiding concern in introducing knowledgeable action is askew from mainstream rationales for caring about the role of knowledge in action. One rationale is driven by an aim to understand knowledge how to A. I do not share this concern, although I will discuss connections and divergences between my view of knowledgeable action and two views on knowledge how in section 8.9.

A different rationale is driven by the thought that knowledge of A-ing contains unusual epistemic features. Knowledge of action is epistemically special. I share an interest in understanding these features. And things will turn out fortuitously. For as I build an account of knowledgeable action,

The Shape of Agency: Control, Action, Skill, Knowledge. Joshua Shepherd, Oxford University Press (2021).
 DOI: 10.1093/oso/9780198866411.003.0008

I will also develop a view on the epistemic specialness of the knowledge embedded in knowledgeable action.

In this chapter I give more attention to knowledge of action than to knowledge of how to A. Until section 8.7 I speak primarily of knowledge of action. At that point it becomes possible to bring knowledge of how to A more directly into play. This is a little bit clunky. Trust me, I know what I'm doing.

8.2 A Word about the Knowledge State

Some philosophers think of knowledge as involving belief. There are different ways knowledge of action could do so. Perhaps in many cases relevant beliefs about what I am doing are very closely associated with my intention (Pavese 2018). Perhaps my intention necessarily involves a belief about what I am doing, or how I am doing it (a la cognitivists about intention, like Marúsič and Schwenkler 2018).

Some philosophers think of knowledge as a *sui generis* cognitive state (Williamson 2002).

Some have argued that knowledge of action could involve (non-belief-involving) intention alone (Campbell 2018).

I am agnostic, in this chapter, about all of this. But I will discuss various philosophers who hold, and speak out of, different views on the kind of state knowledge of action is. It is most natural, I think, to take the view I articulate as a kind of belief-involving, belief-is-distinct-from-intention view. But it could just as well be understood in terms of a *sui generis* knowledge state, or in terms of an intention if intention includes the right kind of belief.

8.3 Knowledge of Action's Epistemic Features

Start with self-knowledge. An agent's knowledge about certain aspects of her mind is in certain ways special or distinct from other sorts of knowledge. Typically, the literature on self-knowledge restricts discussion to self-knowledge of mental states. So, in a recent introduction to the topic, Brie Gertler writes that "In philosophy, 'self-knowledge' standardly refers to knowledge of one's own sensations, thoughts, beliefs, and other mental states" (Gertler 2015: 1). And in a recent book-length treatment of the

recent literature on self-knowledge (Coliva 2016), one finds no discussion of knowledge of action. But arguments are not given for omitting discussion of knowledge of action; the omission is arguably artefactual.[1]

An agent's knowledge of her own action is often said to be special or distinct in much the way self-knowledge of other elements of mind is said to be. Indeed, in my view the consideration of knowledge of action as a part of self-knowledge of aspects of mind more generally proves useful in setting the epistemological task facing the philosopher. In both cases, the task is to clarify what is distinctive about the knowledge in question, to demonstrate that the knowledge really is distinctive in the elucidated way(s), and to elucidate the grounds for claiming that what is distinctive amounts to knowledge.

Move to what might be distinctive about knowledge of action. Elizabeth Anscombe famously asserts that an agent's knowledge of her own action is in some sense non-observational. Anscombe introduces this claim in order to explain the criterion she proposes for distinguishing intentional from non-intentional action.[2] Intentional actions "are the actions to which a certain sense of the question 'Why?' is given application; the sense is of course that in which the answer, if positive, gives a reason for action" (Anscombe 2000: 9). But what is it to give a reason for action? For Anscombe, one cannot give a reason by simply citing a cause of the action. Nor can one give a reason by citing "evidence for supposing the thing will take place" (9). Instead, she suggests we understand the giving of reasons in terms of the acceptance or refusal of the application of the question "Why?" to some bit of behavior.

Two ways of refusing application of this question to one's behavior involve knowledge. First, we refuse application of the question if we answer "I did not know I was doing that" (11). Second, we refuse application if we answer "I knew I was doing that, but only by observation" (14). As Anscombe has it, then, non-observational knowledge of one's action is necessary (though insufficient) for legitimate acceptance of the request for reasons. This is meant to highlight a revelatory truth about the nature of intentional action: unlike our non-intentional behavior, our intentional

[1] In this connection, Lucy O'Brien writes: "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. 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 perception, that have given rise to these debates" (2009: 156).

[2] In explicating Anscombe's view, I need not endorse her distinction between non-intentional and intentional action, nor her way of drawing it.

actions are (normally) those things we intimately know, in the sense that we easily accept and easily answer the request for reasons.

But why does knowledge of behavior drawn from observation render one unable to legitimately accept the request for reasons? In her reconstruction of Anscombe's reasoning on this point, Hanna Pickard (2004) invokes the "intuition that one's knowledge of one's own actions bears some affinity to one's knowledge of one's own mind" (207). This is an interesting thought. As Richard Moran (2004) observes, central to Anscombe's thinking on knowledge of action is the thought that "the agent's own conception of what he is doing is not just another description, side by side with all the others . . . There is a privileged relation, though not incorrigible, between what the agent is doing and what he takes himself to be doing" (Moran 2004: 44). The connection between knowledge's non-observational nature, and an agent's authority in knowing it, follows closely behind this observation. Pickard notes that one's knowledge of an action "is not something about which one is especially authoritative if it is based on observing, as anybody suitably positioned can, what gets done" (Pickard 2004: 207). The idea that we can understand the purpose of Anscombian knowledge without observation as part of an explanation of an agent's authority in knowledge is promising.

What is it, though, for an agent to be authoritative in knowing what she is doing? Alex Byrne (2005) helpfully distinguishes between peculiar epistemic access to an item and privileged epistemic access to an item. Privileged access is understood in terms of an increased likelihood that one's beliefs regarding the item in question will amount to knowledge. And the increase in likelihood is understood as relative to some contrast class. As Byrne's focus is knowledge of one's own mental states, he notes that one could have privileged access to one's own mental states in comparison with one's own access to other items (e.g., beliefs about one's environment), and in comparison with another's access to one's mental states. Peculiar access is understood as "a special way or method of knowing" the item in question (Byrne 2005: 81). One reason for drawing this distinction is that these two types of access are logically distinct. As Byrne notes, some maintain that our access to our own mental states is non-peculiar—fundamentally the same as the access others enjoy—although we may enjoy privilege due to our increased familiarity with ourselves. By contrast, it is possible to maintain that we have peculiar access to our mental states (perhaps via some faculty of inner sense), but that this access is underprivileged, perhaps because "one's 'inner eye' is very unreliable by comparison with one's outer eyes" (81). Another reason for drawing the distinction is that in the good case

there seems to be an explanatory relationship between these two modes of access. As Byrne notes, in good cases "they are connected: the kind of peculiar access that we enjoy presumably explains why we have privileged access" (Byrne 2011: 202, see also O'Brien 2009: 158).

We might apply the same distinction to a difference between one's epistemic access to what one is doing and another's epistemic access to the same thing. On this application, one has peculiar access if one's way of coming to know what one is doing is not available to others. And one has privileged access to what one is doing if one's beliefs about what one is doing are more likely to amount to knowledge than the beliefs of others regarding what one is doing. Ideally, an explanation of one's authority in knowing what one is doing will capture the explanatory relationship between these two modes of access. As I will understand it, then, an agent's authority in knowing what she is doing is constituted by the agent's having both peculiar and privileged access to what she is doing, where the peculiar access helps explain the privileged access.

Thus understood, an agent's authority in knowing what she is doing is in part explained by the knowledge's non-observational nature, because if one's epistemic grounds for this knowledge were fundamentally observational, then it would fail the above explication of authority. If an agent had her knowledge in a purely observational way, she would not have it in a peculiar way. But peculiarity is a necessary condition on my understanding of authority.

Now we ask how an agent comes by this authoritative knowledge. Philosophers have tried various routes. Controversy abounds. In what follows I discuss some of these attempts, and I explain why I find them wanting. Then I offer my own proposal, and explain how it illuminates knowledgeable action as a mode of agentive excellence.

8.4 Intention and its Limitations

Say an agent intends to A, and is in the midst of executing the intention. He is A-ing. In discussing knowledge of A-ing, Kevin Falvey (2000) and Sarah Paul (2009a) appeal to the agent's intention to A to explain his knowledge of A-ing. Their views, while distinct, both appeal as well to the openness of certain act-descriptions (the so-called atelic progressives). Sarah Paul offers an example:

> The truth of the statement 'Jen is crossing the street' does not entail that Jen crossed the street; she might turn around and go back, or she might

> get run over by a truck, without its ceasing to be true that she was crossing the street. (Paul 2009a: 16)

Paul argues that for actions under such descriptions, "if the agent is or has been doing them at all, it is true that he has done them," and thus "the agent can count as so acting even if the goal of the activity is never successfully achieved." For such actions, then, "given that the agent need not ever achieve some terminal endpoint in order to count as doing [such actions]…he also need not have observational evidence that he is succeeding to know he is doing it" (16). Paul claims further that something similar can be true of telic progressives: in their case, "the agent can count as so acting even if the goal of the activity is never successfully achieved. As Anscombe puts the point, a man can be doing something which he nevertheless does not do (as when Jen gets killed halfway through crossing the street)" (16).

How does the agent come by this (non-observational, and therefore at least peculiar) knowledge? For Paul, the agent can do so by inference from his intention,[3] when appropriately situated. Paul elaborates upon the agent's situation as follows. The agent must have favorable background beliefs regarding: a) their ability, b) their circumstances, c) their history as a reliable agent, and d) their "understanding of the way some action descriptions apply partly in virtue of [their] intention in acting" (Paul 2009a: 18).

This certainly seems like one way an agent can come by knowledge of what she is doing. Just as certainly, there seem to be many cases that fail one of these conditions, in which the agent knows what she is doing. Imagine that you have suffered from a recent bout of Miller Fisher Syndrome—a rare reaction to a virus that causes nerve de-myelination, resulting in muscle weakness and disorders of balance and motor coordination, among other things. Say that you have been unable to descend stairs without help for a few weeks, but you awake one morning and decide to try. The prospect of doing so successfully is much less certain for you than it was a few weeks ago. In fact, you may be uncertain whether you can succeed, and you may only have an intention to try to A. (We could run the case with an intention to try to A, or to A.)

[3] While Paul's account knowingly discards some of the epistemic specialness, Falvey's account is rather different. He suggests that knowledge of A-ing consists in knowledge of an intention to A, provided one is entitled to assume the execution of the intention is going well. My worry about this account is similar to my worry about other intention-based views. The scope is too narrow.

What happens when you descend the stairs? Speaking from experience, I can say that one descends with an intentionally heightened awareness of the progress of one's actions. Where once one descended unthinkingly, now one grips the rail, monitoring footfall, muscle tension, balance, and speed. One is very aware of what one is doing in part because one is aware of the potential for difficulties. One is poised to observe mistakes, to make corrections. (This may happen to healthy agents in mundane cases as well, such as descending stairs in traffic, or descending stairs the height of which is unfamiliar.) It may happily be that no difficulties are encountered. One descends the steps exactly as planned. As one does, one knows what one is doing. One not only knows that one is trying to descend the stairs. One knows that one is descending the stairs. An inference based upon intention plus favorable background beliefs does not seem to capture this.[4]

None of this is to deny that intention, given the right background conditions, can be an important epistemic resource for agents, and can generate knowledge of action. My contention is that intention cannot be the only

[4] Kieran Setiya also argues that agents come by their knowledge of A-ing via their intentions. But Setiya does not lean on inference. And he emphasizes slightly different background conditions. For Setiya (2008), agents must know how to A, and must know they are able to A. These features justify a decision—the formation of an intention to A, and a corresponding belief. How might this get us all the way to knowledge of action?

In response to the trenchant criticisms of Sarah Paul (2009b), Setiya (2009) considers a case in which an agent is *Φ*-ing by performing basic intentional actions ABC. Setiya claims the agent knows she is *Φ*-ing by knowing she is ABC-ing. So we have to ask how an agent can (for Setiya: non-inferentially, non-observationally) know she is ABC-ing. Setiya says this:

> How do I know that I am doing ABC? We can say at least this. If I know how to take those basic means, this knowledge consists in the disposition to execute the corresponding intentions. Since I have this disposition, it is no accident that, when I intend and thus believe that I am doing ABC, I am doing so in fact. Exercising basic knowledge how ensures non-accidentally true belief. To say this is not to endorse an epistemology on which its being no accident that a belief is true suffices for it to count as knowledge. But it does pre-empt a residual source of skepticism, that the truth of beliefs formed without sufficient prior evidence could only be a matter of luck. When they are constituted by intentions and one knows how to perform the relevant actions, that is not the case. (Setiya 2009: 136)

When I know how to ABC, the belief that I am ABC-ing can be non-accidentally true because the knowledge how disposes me to execute the intentions that constitute the relevant beliefs.

Against accounts like Setiya's, Schwenkler argues that "it is hard to see how [such accounts] can work in cases where failure is quite likely, e.g. ones involving especially hard-to-perform actions, challenging conditions, extremely weak-willed agents who tend not to follow through on their intentions, and so on" (Schwenkler 2012: 733). Since agents who act successfully in difficult circumstances do often have knowledge of their actions, it looks like Setiya's account leaves out important cases in which we know what we are doing.

epistemic resource available to acting agents. Nor can it undergird the account of knowledgeable action I introduce below.

In any case intention-based accounts are in a way unnatural. They have difficulty capturing our non-observational knowledge of what is actually happening—of how the action is actually unfolding. The agent who knows she is A-ing by intention alone remains very much in the dark.[5]

8.5 Perception and its Limitations

If intention-based accounts leave agents in the dark, it is natural to wonder if perception provides a needed supplement. Most philosophers think not. Perception, they think, can be at best an enabling condition, and not a part of any justification, for knowledge of action (Falvey 2000, 32). John Schwenkler (2011, 2012, 2015) disagrees, claiming the following:

> There seem to be cases where the role of perception in agent's knowledge is plainly justificatory, as evidenced by the way perception would be cited in answer to a question of the form "How do you know that you are X-ing (say, replenishing the water supply of the house)?"—"Because I can see that I am," the agent might reply, thereby identifying perception as the justificatory ground of his or her knowledge, and not just something that enables him or her to know this thing non-empirically. (2015: 26, fn. 63)

To understand more fully Schwenkler's position, I should bring out an element of his broader picture. Schwenkler grounds the formal causal

[5] The same criticism applies to Thor Grünbaum's (2013) interesting discussion of a role for perceptual demonstratives in the knowledge of action. This is because perception's role, for Grünbaum, is to help specify the content of an intention:

> Usually when acting manipulatively on objects in our environment, intentions which make general or descriptive reference to objects of action are often not specific enough to initiate and guide our actions. Intentions to throw the paper in the waste basket or to write one's signature on some form are not specific enough. The agent needs to know which particular object is the basket or form in question before she can initiate her action. We can think of perceptual attention as having the role of selecting the right object in the world and thereby specifying the prior intention. The role of attention is to transform the prior intention (too unspecific to initiate and guide action) into a demonstrative intention: the intention to φ with this o. (Grünbaum 2013: 301)

role of knowledge emphasized in some interpretations of Anscombe (e.g., Moran 2004) in its efficient causal role. Consider the following passage:

> To explain what makes agential knowledge practical it is not enough just to say that this knowledge is 'what gives the descriptions under which what is going on is the execution of an intention' (Anscombe 2000: 87)—or rather, we seem not to have explained how our movements ever could be so described, until we explain how the knowledge of what one is intentionally doing makes a difference to the character of these movements themselves, and is therefore a knowledge of oneself as an agent, and not a merely passive subject of motion. (2015: 9)

In setting out his account of practical knowledge, Schwenkler is also diagnosing a flaw in Anscombe's. While sufficiently powerful agents need not check in on the progress of their actions, human agents can do very little without perception. Anscombe is aware of this, but argues that perception is often a mere aid for action execution, not an indispensable basis for knowledge of action. But, as the failures of intention-based accounts make clear, it does not look like perception is a mere aid. Knowledge of what a human agent is actually doing appears to require information that could only be gleaned from perception.

It is at this place that the grounding of knowledge's formal causal role in its efficient causal role does important work for Schwenkler. He wants to explain how knowledge of action could be genuinely practical and non-observational while being grounded in perception. He claims that in many cases "perception may serve as an 'aid' to action just by informing the agent of whether she is or is not acting as she means to: its guiding role is inseparable from its grounding role as a source of agential knowledge" (2015: 26). Ultimately, then, for Schwenkler knowledge grounded in perception "is not a mere 'monitoring' of an independent event, but an integral part of the process by which a person brings about the very action that it is a representation of" (2015: 27).

I like much in this view, but I have a few worries. First, such language suggests that Schwenkler thinks that perception represents what is happening as an action. It is debatable whether it does.

Second, even if perception sometimes represents bodily movements or other events as action, it is unclear that it does so in a way that could underwrite an agent's authority in knowledge of action, or an agent's claim to have non-observational knowledge of what she is doing. For if perception can

represent bodily movements as action, presumably it can represent your as well as my movements as action. If perception is the sole ground of non-observational knowledge of action, how is my knowledge of what I am doing distinct from my knowledge of what you are doing? The worry is that this account jettisons rather than explains knowledge's independence from observation, or an agent's authority in this knowledge.

Third, not all information drawn from perception could play a formal causal role in determining what one is doing, even if all information drawn from perception could, in principle, play an efficient causal role in influencing what one does. For consider the following four roles for perception in action. First, some information drawn from perception signals that one is succeeding in action. One kicks a football, feels the impact at just the right spot on one's boot, sees the ball take off at the right angle. One is doing as one intends. Second, some information drawn from perception signals that one is making a mistake in action. One kicks a football but feels one's leg out of place in the swing, hears something odd in the way the boot strikes the ball. Third, some information drawn from perception signals unpredicted effects of one's action. One kicks a football and notices a bystander watching its flight. Fourth, some information drawn from perception signals features of the world not immediately relevant to one's action. One kicks a football and notices storm clouds on the horizon.

While the first two types of information might play a formal role with respect to an experience of what one is doing—either succeeding or not (perfectly) succeeding—the second two types of information may play no such role. Certainly as one acts one continues to experience the world, and not every part of the world is relevant to what one is doing or to one's knowledge of what one is doing. The latter two kinds of information may in principle, however, contribute to how one acts, even though neither is represented or experienced as something the agent intentionally does.

Fourth, no one should deny that perception can add shape to action. What the agent does, and how she does it, cannot be fully explained in many cases without appeal to perception, and to perception-action and perception-cognition links. But what does knowledge have to do with this? Schwenkler wants to explain "how the knowledge of what one is intentionally doing makes a difference to the character of these movements themselves" (2015: 9). But the knowledge here seems superfluous. Perception does all of the work.

I return to this last worry below, where I argue, in effect, that perception can be important for knowledge of action—as well as for knowledge of how

to act, and for knowledgeable action—when and because the agent puts the relevant information to use in a certain way.

8.6 Conscious Awareness and its Limitations

8.6.1 Pickard and Bodily Experience

Hanna Pickard (2004) is motivated in part to explain what she calls an "overlooked feature of our knowledge" of action. She explains this feature as follows:

> Actions are events: they occur during periods of time. Knowledge of such events must be sensitive to their course through time. Part of what is distinctive about one's knowledge of one's own actions is the nature of this sensitivity. Not only does one normally believe that one is acting if one is, but if one believes that one is acting, then normally one is. One's knowledge keeps near perfect time with the actions themselves.
>
> (Pickard 2004: 206)

I agree. When acting we do not experience ourselves as acting under a loose description that involves only our intention and related beliefs. We experience ourselves acting—doing things in the world. Pickard is sensitive to this important phenomenon:

> Certainly we do seem to experience ourselves acting. If one is building something, it is not just that one knows what one intends, and can see what in fact results: intuitively, one can be aware of doing the building.
>
> (Pickard 2004: 212)

How best to capture this phenomenon? Pickard's proposal adverts to the content of bodily experience. For Pickard, "one can be aware of acting in virtue of being aware of one's body—that with which one acts—from the inside" (Pickard 2004: 212).

The claim is not that bodily experience contributes to one's awareness of acting, but rather that it grounds one's awareness of acting. Is this true? Certainly Anscombe seemed to reject such a view. For Anscombe, the sensations one has in one's body—sensations of pressure at various body points, or of stretching in various muscles, for example—are not sufficient

to explain the non-observational knowledge one has of the position of one's body parts. Anscombe seems to ignore the possibility that bodily experience could have higher-level content than such sensations, however. Pickard seeks to capitalize on this possibility. She asserts that not only does bodily experience have high-level content, it has content that describes the action one is performing at a time:

> We hear not just sounds but words; we feel not just an impenetrable surface, but a table. Given that body awareness is a form of perception, we should be willing to consider the possibility that its content is comparably rich: when one acts with one's body so as to pick up a cup, the content of the awareness one has of one's body is as of picking up a cup.
>
> (Pickard 2004: 220)

In order for bodily experience to ground knowledge of action, Pickard not only needs bodily experience to have rich (or so-called "high-level") contents (see Bayne 2009). She needs the contents to ground beliefs about action rather than about mere bodily movement. It would be more straightforward to understand how this happens, and more in line with Pickard's intentions here, if bodily experience presents the body as acting. Pickard makes the following claim: "one can be aware from the inside not just of one's arm rising, but of raising it" (Pickard 2004: 206). Later in the same paper she claims that "the awareness one has of one's body as acting not only facilitates the practical execution of action. It can also serve to sustain one's present-tensed knowledge that one is acting over the course of the action" (227).

Does bodily experience present the body as acting rather than as moving? This claim is difficult to assess. One reason for this is that bodily experience is typically accompanied by a range of other experiences, both perceptual and non-perceptual. One rarely if ever has an experience that only involves the elements that make up bodily experience. So it is difficult to find or to develop the kind of contrast that we need.

One might think that anarchic hand syndrome (AHS) could shed some light on Pickard's claim. Agents with AHS perform actions with an anarchic hand while denying that they are the driving force of the actions (Marchetti and Della Sala 1998). These agents still experience their hand as moving, but they do not experience its movements as their actions. One might think, then, that such cases are a problem for Pickard. But it remains possible that these agents experience their anarchic hand as acting, and that they distance themselves from the action for other reasons.

Even if this is right, however, the fact that anarchic hand patients commonly describe themselves as surprised by the actions of their anarchic hand indicates that the bodily experience accompanying the hand's movements is not sufficient for non-observational knowledge of action. In the agentive phenomenology literature, it is often claimed that the conscious performance of an action carries with it a sense of mineness. The way Horgan, Tienson, and Graham (2003) explicate this aspect, "You experience your arm, hand, and fingers as being moved by you yourself, rather than experiencing their motion either as fortuitously moving just as you want them to move or else as being transeuntly caused by your own mental states" (329). One problem raised by AHS is whether bodily experience is alone sufficient to capture this important aspect. Although forms of AHS can have sensory consequences (e.g., AHS that stems from damage to the parietal lobe (Jenkinson et al. 2015)), AHS is not specifically a deficit of bodily experience. This indicates that bodily experience is not sufficient for mineness phenomenology. If we wish to capture non-observational, practical, authoritative knowledge of action, bodily experience must at the very least be supplemented or integrated into other grounds.

8.6.2 Peacocke and Awareness of A-ing

Christopher Peacocke (1999, 2003) advances an account of knowledge of action on which an agent's conscious awareness of A-ing provides a reason for judging that she is A-ing.

Peacocke argues that conscious awareness has the right kind of content and the right kind of connection to an agent's judgments to provide the required link to knowledge. For when an agent takes elements of her conscious experience at face value, and her conscious experience is related to the world in the right way, the judgment that results is not only truth-conducive. In addition, it can "be philosophically explained in terms of the nature of the intentional contents and states involved in the transition" (2003: 52). Reliability or truth-conduciveness alone is not sufficient. Reliability must come about in virtue of the "nature of the intentional contents and states involved" (52).

I do not wish to disagree here with Peacocke's epistemic proposal. My worry concerns his account of the awareness of A-ing. I have criticized this account in an earlier paper (Shepherd 2017b). Here is the gist.

First, the experience of bodily action seems, phenomenologically, to involve perception to some extent. When we experience ourselves moving

our bodies, a feature of the experience is that the body is moving. We need not deny that illusions of movement can be caused by tryings in some cases. But it is implausible that a trying alone can account for the actional experience of the body moving.[6]

Second, Peacocke appeals to work by Anthony Marcel and colleagues to support his view (see Marcel 2003). I argue that the work does not, in truth, offer the required support. My earlier explanation is, I hope, sufficient:

> Marcel's basic result is this. Using vibro-tactile stimulation to the elbow, it is possible to put agents in a position in which they falsely experience their arm in one location and so intend to move their arm in one direction in order to accomplish a goal. But due to the illusion induced by stimulation, these agents must actually move their arm in the other direction to accomplish the goal. It turns out that agents are accurate in moving their arms in the correct direction. But 60–70 percent of agents in Marcel's study show no awareness of having moved in a direction opposite to their consciously reported intention. From this Peacocke concludes 'The content of their trying (or some event causally related to it) seems to cause the content of their impression of action, even though the actual motor instruction issued requires, and produces, movement in the opposite direction' (Peacocke 2003: 122). But it is not clear how strong this evidence is. For, as Marcel notes, it is possible that these agents had no awareness of the specific character of their arm movements: their strong belief that they were moving their arm in one direction might have overridden whatever arm-related experience

[6] I would also apply phenomenological complaints to Lucy O'Brien's (2009) account of knowledge of A-ing. Like Peacocke, O'Brien draws a link between awareness and knowledge. Her account of awareness is different. For O'Brien the awareness of A-ing is not caused by a trying to A. Rather, to consciously act is to act in "a certain way or mode"—a conscious way or mode. And this is to act "on the basis of an evaluation of the possibilities open to us" (O'Brien 2009: 184). When an agent acts on the basis of an evaluation of action possibilities, O'Brien claims she acts with a sense of control. O'Brien further claims that to act with a sense of control is to act with knowledge-enabling "agent's awareness":

> I act consciously when I am agent aware of my action, and…I am agent aware of my action when it is something I actively control. Our actions are those things we know, not by observing them, or by reflecting about them, or accepting some presentation of them, but rather by actively engaging in them. (O'Brien 2009: 183)

I do think agents sometimes act with knowledge-enabling agent's awareness. This is what was going on, I think, in my earlier stair-descent case. My chief problem with O'Brien's account is that it seems the wrong description of the contents of the experience of acting. The evaluation of action-possibilities takes place prior to the enacting of whatever possibility one settles upon. So if we want to explain our active engagement in the action, we need to say something about the content of the experience as we engage.

> was present, leading to their false reports of arm movement. Second, as Marcel also notes, 30–40 percent of participants did have conscious access to the conflict between intention and arm movement. Marcel reports that 'All of the latter subjects commented in one or another way that there was something peculiar, or that while they had intended to move in one direction they had the impression that they had moved in the other' (2003: 166). That the relationship between trying to A and having an experience of A-ing Peacocke asserts is not universal—even if it is present in most of the 60–70 percent of subjects in Marcel's study, a claim I have questioned—casts doubt on this as a general account of the experience of acting.
>
> (Shepherd 2017b: 428)

A further problem is that temporarily paralyzed agents can experience tryings without experiencing any associated bodily movements (Shepherd 2016). So it seems that this account of the experience of acting will not work.

8.6.3 My Suggestion and its Limitations

The account I offer of the experience of acting's contents appeals to work on informational integration in the etiology of our conscious experience. I argue that just as some of our perceptual experiences result from the integration of contents in separate modalities (see O'Callaghan 2015), such as vision and audition, our experiences of acting result from the integration of contents due to intention (and perhaps motor command) as well as various perceptual modalities.

I claim:

> The experience of acting's distinct unity stems at least in part from a sub-personal process of construction that integrates spatial and temporal contents drawn from agentive and perceptual processing and attributes the result to the same unfolding event: the action.
>
> (Shepherd 2017b: 443–4)

The upshot of this account—the "distinct unity" claim—is that the experience of acting is legitimately of acting. When consciously acting the agent is experiencing doing a thing—it is probably more accurate to say that she is experientially doing a thing—in virtue of the contents drawn from

intention and motor command, and the agent is experiencing doing a thing, in virtue of the information gleaned from perception.

One might run an epistemic angle on this proposal as Peacocke does. Or one could try out a kind of dogmatism (a la Pryor 2000), as I suggest in (Shepherd 2017b). This can, I believe, get you to knowledge of action in some cases—the cases that involve this kind of experience of acting. So we now have at least two sources that may very well generate epistemically special knowledge of action—intention and the experience of acting.

But this is not the full story. For not every case of action involves the experience of acting. Sometimes this is because the kind of action at issue does not involve the right kind of intention-perception links. Sometimes it is because consciousness is not involved in the action in the right kind of way. So the knowledge one gets via this route is, again, restricted.

And there is a further issue, somewhat separate from the epistemic account of the knowledge of action. I am ultimately interested in knowledgeable action. And the role of knowledge in action is broader than these sources could indicate.

Notice: the knowledge acquired via intention or the experience of acting seems otiose for the purposes of action execution. The relevant work is done by states and processes—intention, perception, experiences—that do not amount to knowledge. What would be good, I argue presently, is if an agent could use experiences of acting, along with other resources at her disposal, to come by a knowledge that contributes to her exercise of control in acting.

8.7 Practical Reasoning in Action

Here's the big secret. An agent has knowledge of action that she can use—knowledge of A that contributes to the A-ing, as well as knowledge of how to A—when she is poised to engage in practical reasoning in action.[7] This covers knowledge of A-ing, as well as knowledge of how to A.

Practical reasoning in action is the familiar notion of practical reasoning, applied to and during action. I do not endorse a specific account of practical

[7] Rödl (2011) develops an Anscombian view of intentional action that depends upon claims about practical reasoning and knowledge of action. There are some affinities of form between his claims and mine, but important differences exist. Mine is not an account of intentional action, but of intentional action that uniquely involves knowledge. Further, my account is liberal about the structure of practical reasoning—no particular claims about the nature or form of the premises or conclusion of practical reasoning are assumed.

reasoning, or of reasoning, here. In general, practical reasoning is a response to uncertainty about some element of the action: about whether to behave in a certain way, whether to continue one's plan or modify or abort, or about whether one can succeed exactly as planned (Camp 2009, Shepherd 2015a). It is then a set of psychological processes directed towards figuring out whether one is succeeding, or whether one can succeed, or what the right answers to one's questions about whether and how to act might be.

So, as one slowly descends the stairs, one's heel may clip a stair leading to a lurching feeling. This may lead to an automatic reaction, to be sure. But in some cases, it will lead to a rapid form of reasoning. One will envision the possibility of gripping the rail harder, and then envision the possibility of halting downward progress by leaning back or sitting down, and then one may envision the possibility of skipping the next step to better land on the step just beyond. Very rapidly, one is calibrating levels of confidence in these possibilities, and the risk of failing at any of them. One may judge that the best move is the least risky, and decide to sit down. As Marúsič and Schwenkler (2018) put it, "practical reasoning is the reasoning by which an agent determines what she is going to do, insofar as what she will do is up to her. That is, insofar as matters are up to the agent, she is in a position to determine how they will unfold, and she does this through her practical reasoning" (310).

It is important to note that I do not require that practical reasoning be actually engaged—that the agent be actively figuring out how to act. I require that the agent be poised to engage in practical reasoning. The poise I have in mind entails activated beliefs about what the agent is doing. This is because in some cases, agents anticipate trouble. Or perhaps agents simply wish to act in a way that involves heightened awareness (where this awareness entails poise to engage in practical reasoning). In such cases agents act with poise. This is a state of readiness that involves awareness of what is going on related to the action, often leading to pick-up of knowledge via various available sources where the pick-up is a part of a disposition to use what is picked up. Poise can be triggered by registration of difficulty, leading to practical reasoning. But it can be in place beforehand as well, before practical reasoning is actually engaged.

The set of psychological processes that qualify as practical reasoning are, in my view, diverse. Some views of reasoning that emphasize attitudes towards complex propositions "in which propositions may themselves figure as objects of which properties and relations are predicated" (Pettit 2007: 498). Such attitudes may play important roles in some circumstances. Tyler Burge speaks of agent's deploying higher-order attitudes in order to "take

rational control of one's reasoning" (Burge 2013: 74). These are useful tricks for limited agents, like human agents. Given limitations in knowledge and functioning we sometimes need to explicitly review our first-order reasoning. Of course there is rarely time to do so. So in many cases our practical reasoning does not take this form.

I noted in chapter 2 that the philosophy of cognitive science identifies a number of ways an agent's practical reasoning may involve (without, perhaps, being fully exhausted by) non-propositionally structured representational states: map-like representations (Camp 2007), or analogue magnitude representations (Beck 2014), or mental imagery (Gauker 2011), or combinations of these (Shepherd 2018a).[8] In other work I have argued that while reasoning via such states may limit agents in some ways, there may be important benefits as well. In skill development agents navigate a tension between efficiency and flexibility—automatization increases efficiency (in the good case), while decreasing flexibility and sensitivity to reasons (Shepherd forthcoming-c). Something similar is true of reasoning via non-propositionally structured states. The types of transitions afforded by map-like or imagistic states may be limited. But the manner of transitions afforded may be just what the agent needs. Some formats are, in some circumstances, or embedded in some systems, more usable than propositionally structured states (Camp 2007; Shepherd 2018a).

Consider the use of what Robert Briscoe (2008) calls "make-perceive." Agents often voluntarily or intentionally generate mental imagery in order to augment the current experiential field or to adumbrate nearby possibilities. This imagery may be motoric or involve various perceptual modalities, depending on the needs of the agent. Developing a facility with this capacity is extraordinarily helpful for the smooth control of action. Briscoe (2018) considers the example of a rock-climber. When climbing a wall, a climber needs to execute several sub-actions that balance accuracy of reach, distribution of weight, strength of grip, position of the feet, accessibility and climbability of upcoming patches of rock, and so on. Skilled rock climbers rely on enhanced perceptual recognition of elements on the rock wall, and enhanced abilities to deploy motor and perceptual imagery as they work through the best ways to execute the action of climbing the wall (see Pezzulo et al. 2010). Plausibly, doing so involves assessments of relations of reason

[8] I suppose these states may relate the agent to propositions, without themselves being propositionally structured. Whether one cares will likely depend upon how invested one is in certain proposition-involving construals of knowledge how.

between ways of executing the action, goals, one's own abilities, and available physical resources.

Notice that the involvement of knowledge in these instances of reasoning may depend upon context- or domain-specific forms of control or skill. Katalin Farkas (2017) has recently emphasized just how context-dependent much of the knowledge driving action may be. In particular, Farkas argues that many of the answers to embedded questions that constitute "knowledge-wh" may be available only "through, or with the aid of, a complex sensory experience" (863). (A quick aside: it is fine to think of my "knowledge of how to A" as a kind of knowledge-wh—knowledge of what to do, when to do it, and so on, are ways of knowing how to A.) To illustrate:

> When the sailor knows she can lean out this far, the content involves a proprioceptive presentation. If I pointed at her and said 'one can lean out that much,' I would not manifest the relevant knowledge-wh, because the content of my knowledge would involve a visual, third person, rather than a proprioceptive presentation. It seems very likely that the ability to perform the maneuver, at least in some cases, requires a proprioceptive presentation. If so, then my knowledge may be sufficient for some purpose, but would not be sufficient for the purpose of performing the maneuver. The reason is that the extent one can lean out would not be presented under the right mode of presentation. (Farkas 2017: 864)

In applying Farkas's insights to knowledge of (bodily) action, I would emphasize an addition. First, it is plausible that much of the knowledge relevant to action control relies on content given not just in sensory, but in sensorimotor experiences—experiences of acting. So we highlight an additional way that conscious experience is relevant to the knowledge of action, even though it cannot be the full story. The conscious experiences one has while acting not only furnish one with materials relevant to one's activity of practical reasoning. In addition, epistemically relevant properties of the action itself are given in a way that only the relevant experience can provide.

Now, practical reasoning in action embeds and depends upon judgments about what one is doing, and how one is doing it. One engages in practical reasoning usually because uncertainty arises about whether what one is doing will lead to success, or whether one's plan for how to do so will lead to success. The engagement of practical reasoning in action gains structure from the action one is executing. In such reasoning, one is seeking to answer questions about whether to proceed as planned (as intended), or about how best to change the plan in light of extenuating circumstances. Sometimes

one is engaging in a form of practical reasoning that is more like monitoring, or waiting: the agent is asking whether she will pull off the next step as planned, and she is poised to intervene in case of a mistake. In all of this various judgments about what one is doing, and how, are more or less taken for granted. But they are there. They set the background against which one reasons. One is reasoning about the best way to A, or about whether to continue A-ing. Such reasoning embeds judgments that one is A-ing, and often embeds judgments about how the A-ing is going at a time, at a relatively fine grain of detail.

These judgments about what one is doing, about how one's action is progressing, may be informed by a range of sources—one's intention, elements of perception, one's experience of acting. But the already-discussed accounts of the epistemic credentials of knowledge via these sources do not transfer. When one acquires knowledge of what one is doing due to a process of practical reasoning, the epistemic account is importantly distinct.

8.7.1 The Epistemic Account

When an agent is engaged in practical reasoning, she is making a series of judgments. These judgments form beliefs. We care in this instance about the beliefs about what the agent is doing, and about how the agent is doing it or will be doing it. These beliefs amount to knowledge when, and because, the agent has a sufficient degree of control over her practical reasoning.

What is it to have control over practical reasoning? The basic features of the account of control developed in chapter two apply. But there I was concerned with control as indexed to an agent's plans, or to her goals. By now we have a sense of the broader spaces in which agents move. So we understand, for example, that agents can be assessed according to goals that they do not currently possess—goals proprietary to action domains. A more expansive anchor for assessing control over practical reasoning is possible.

An agent has a sufficient degree of control over her practical reasoning in action when her reasoning, in the relevant circumstances, reliably finds ways to successfully A, or, if the circumstances are genuinely not favorable to A-ing, finds the best available options (given background goals and preferences and additional goals accessible to the agent). Often action circumstances afford multiple actions, and sometimes practical reasoning suggests a way to satisfy background goals or desires, or to promote background values, by updating or changing existing plans. Agents may reason their way to very different courses of action, but this may plausibly be considered a

success if the result is preferable to the possible world in which the agent did not engage in reasoning, given the agent's pre-existing plans, preferences, and commitments, or given the action domains in which she is currently engaged.

I said the agent's beliefs about what she is doing and how amount to knowledge when, and because, the beliefs result from a process of sufficiently controlled practical reasoning. Whence comes the because?

When an agent has sufficient degree of control with respect to her practical reasoning in action, in the relevant circumstances and the action domain they fall under, and when she is exercising this control, then she has entitlement to the judgments she makes while conducting that process.

Entitlement is a kind of warrant, an epistemic status a cognitive state can possess that, given the right additional features (about which I'm here going to be neutral), can contribute to that state's qualifying as knowledge.[9] Entitlement differs from justification—another type of warrant—in the following way. While justification requires a subject to be able to articulate the reasons she has for some belief or judgment, entitlement does not require a subject to be able to articulate or conceptualize her reasons for a belief or judgment. All that is required for entitlement is that a process that confers the status of entitlement provides a good route to the truth of the entitled state. So, for example, Burge writes that our entitlement should be "connected to the truth of our judgments about our thoughts" in a non-accidental way (1996: 103). And Peacocke writes that an account of entitlement must show how it is "reasonable to employ the ways of reaching beliefs that the theory describes as entitling, and those ways must be appropriately connected with the truth of the beliefs reached in those ways" (1996: 117–18).

The way I have set this up, we have a good case and a bad case. In the good case, the agent has sufficient control over her behavior, and she exercises sufficient control, by way of having and exercising sufficient control with respect to practical reasoning. In the good case the agent has entitlement, and because of that, her judgments amount to knowledge. The agent only has entitlement in the good case.[10]

[9] So, Tyler Burge says of warrant that "Being epistemically warranted in having a belief is having the belief in a way that is good for having true beliefs, given limitations on one's information and cognitive capacities…if one is warranted, one's belief is held through a natural competence that is epistemically good—conducive to the belief's being true and, usually, to the belief's constituting knowledge" (2013: 489).

[10] Readers might note similarities between my account of entitlement and Burge's account of entitlement for self-knowledge of attitudes like belief, desire, and intention. The accounts are similar in some ways, but distinct in others. Burge's applies primarily to knowledge of the

In the bad case, the agent lacks sufficient control over behavior. (This entails that she will not exercise sufficient control.) This may be because he lacks sufficient control regarding what he intends to do. Or the problem could lie somewhere in his cognitive apparatus. Either way, practical reasoning in those circumstances would not be a good route to finding ways to A. Nor would the agent's judgments about what they are doing amount to knowledge.[11]

For example, if you were to place me on the television show *Dancing with the Stars*, I would be very uncertain about how to perform. In my case, practical reasoning would not be a rational response to uncertainty. For I am so bad at formal dancing, and I know so little about it, that my practical reasoning is unlikely to turn up any solution that could guide me away from uncertainty and towards success.[12] At this low level of control, I am unlikely to judge correctly regarding what I am (intentionally) doing. For I am doing little except for flailing about, and that unintentionally.

8.8 Knowledgeable Action

Knowledgeable action is a subset of intentional action. It is intentional action guided by knowledge of A-ing, and knowledge of how to A, where the knowledge states qualify as knowledge in virtue of the agent's entitlement to her judgments, the entitlement comes about in virtue of the agent's

attitudes; mine to knowledge of action as well as to knowledge of ways of acting. Burge rejects the thought that the entitlement obtains in virtue of "the reliability of some causal-perceptual relation between cognition and its object" (1996: 98). He emphasizes instead the role of an agent's judgments in critical reasoning, and a constitutive relation between "the judgments about one's thoughts and the judgments being true" (98). I actually emphasize the possibility of such a constitutive relation (see section 8.8.1). But it obtains in virtue of control, which is understood, broadly, as involving a reliable causal relationship between the agent's practical reasoning and her behavior. So I doubt Burge would accept my account, even though I think it could be used—via the account of control—to illuminate self-knowledge of the attitudes as well. For Burge's argument relies on the idea that critical reasoning, which involves knowledge, is a kind of rational control over reasoning. And one could understand rational control as control over reasoning that, in virtue of the entitlement the control provides, involves knowledge of relevant reasons.

[11] At least, not because of the entitlement I am here articulating. Perhaps the agent could infer from intention or perception what she is up to—likely, failing, or only trying to A—but in such a case, the knowledge of trying to A would not come from the agents control regarding practical reasoning.

[12] I would need a good teacher, and the absence of an audience—maybe under a bridge somewhere—before I made any progress.

control over practical reasoning plus behavior, and the control over practical reasoning plus behavior explains the fact that the knowledge guides the agent's action.

8.8.1 The Radical Case

This account allows for a radical case in which the agent becomes uncertain about her ability to succeed, begins to engage in practical reasoning, makes a judgment that is disconnected from her prior evidence about what she was doing or trying to do, and so now holds a belief that she is A-ing in way X (when before she was A-ing in way Y, or B-ing in way Z). The case is radical because it allows that if the agent has control over her A-ing in way X in virtue of having control over the practical reasoning that can direct her A-ing in way X, she now knows that she is A-ing, and she knows that X is a way to do so, and the way she is doing so. And this in spite of whatever she was doing before the judgment.

So it goes. I think in such a case the agent redirects herself—when she judges she is A-ing in way X, that's what she is then doing, in virtue of the judgment. Such a judgment need not be a part of her intention. Instead, what happens is that the judgment causes a relevant intention to A in way X. Of course the connection is not foolproof. But control over practical reasoning is in part constituted by these kinds of connections. In the good case—the case in which the agent exercises sufficient control—such a judgment causes a corresponding intention. If such a link fails, the agent is not exercising sufficient control over practical reasoning, and the entitlement disappears.

8.9 Knowledgeable Action and Related Proposals

I wish to draw some connections, and note some departures, from related work that has helped me to see the relevant issues more clearly.[13]

Imogen Dickie (2015, building upon her 2012) offers an account of practical knowledge. She begins with an example evocative of the problem

[13] I came too late upon Kim Frost's (2019) view (based upon a reading of Anscombe) that practical knowledge can be understood as a kind of capacity. This view seems to bake control into exercises of the capacity for practical knowledge in a way that I find interesting. I hope to address this kind of view in later work.

facing intention-based accounts of knowledge of action. Robin Hood forms an intention to hit a target some distance off. Does Robin know that he is going to hit the target? We could extend this to the timespan of action—as he is shooting, does Robin know that he is hitting the target? Dickie says yes, and her explanation reverses the traditional order of explanation from truthmaker to knowledge-qualifying state. (Dickie speaks of the relevant state as both a belief and the cognitive commitment component of an intention.) She comments:

> In this case, the status of the belief as more than just luckily true is not secured by any role the truthmaker plays in the belief's formation. It is secured by the role the belief plays in bringing about the truthmaker. Given Robin's skill, when he commits himself to hitting the target, it is not just a matter of luck that this commitment ends up fulfilled… Rather, Robin's skill enables his intention to be a non-lucky generator of its own fulfilment… So whether an intention counts as knowledge does not depend on how it is formed. It depends on whether it plays an appropriate role in generating its fulfilment. (Dickie 2015: 89)

Dickie is aware that this account depends upon some account of the "appropriate role." She doesn't wish to commit to one. However, given similarities between her account of skill, and my account of control,[14] I would offer my account of non-deviant causation as a useful strut. (I don't know whether Dickie would accept this.)

I like Dickie's idea of appealing to the agent's control (Dickie says skill) to explain why some relevant state qualifies as knowledge. I have used this idea, but with a different target. Dickie's concern here is the attitude of cognitive commitment to some plan. That's either an intention, or a component of some intention. Such a state may qualify as knowledge (I'm neutral), but even if so, it's role as knowledge seems secondary to its role in action execution. For Dickie, controlled action manifests knowledge. But the knowledge does not do work in explaining the agent's success—it is rather explained by her capacity to succeed. (This point leads Habgood-Coote to complain that

[14] Here is how Dickie states things in her 2012: "S is a skilled Φ-er iff, for appropriate range $\sum$ of situations r, in most cases, if S were to intend to Φ in r, S's intention would lead S to act in some [way] $w \in f(r)$, where f is a function taking each r to the set of reliable ways for S to Φ in r" (Dickie 2012: 739) What Dickie is calling skill I would call control. There are slight differences, and potential disagreements in how we would fill out "appropriate range of situations," and in how we would understand the role of intention.

"Dickie's view makes knowledge-how into an epiphenomenal by-product of skill" (2019: 14).)

I am interested in a broader range of states. When engaged in practical reasoning, the agent's action not only expresses knowledge (if it does), it embeds knowledge of action and knowledge of how to act, which then plays a role in the control of that very action. That's one thing I wish to capture. Furthermore, when engaged in practical reasoning, the agent must not simply look for the best ways to execute her intention. They must also remain open to radical changes of plan—changing or dropping intentions, for example. Such may be the demands of practical rationality, or the force of the agent's reasons, in a situation. I also wish to capture this additional layer of possibility. This is why we need a broader notion of controlled practical reasoning, as opposed simply to a notion of control at executing a plan-state.

I apply Dickie's idea that the explanation of the knowledge involved in action runs through the agent's control. In my case the knowledge involved is broader, and the control involved is different.

I turn to a recent discussion of knowledge how by Joshua Habgood-Coote. Habgood-Coote offers an account of knowledge how called the interrogative capacity view:

> For any context c, subject S, and activity V, an utterance of 'S knows how to V' (in its practical-knowledge ascribing sense) is true in c iff c has associated with it a set of practically relevant situations {F1, F2, . . . }, and, for all (or at least most) Fi that are members of {F1, F2, . . . }, S has the capacity to activate knowledge of a fine-grained answer to the question, how to V in Fi?, in the process of V-ing. (Habgood-Coote 2019: 92)

The capacity to activate knowledge is, for Habgood-Coote, in a way disjunctive. An agent might activate knowledge by learning new information, or she might exercise pre-existing knowledge. And there is an additional layer of complexity. For the agent might activate knowledge of how to V in the process of V-ing, Habgood-Coote states, simply by V-ing in some specific way. We are told that "Merely being able to think an answer to a question is not enough for knowing how" (91). The knowledge and the action are intertwined. "The ability to answer the question how to V? on the fly is at the same time an ability to answer the question how to V? by doing V, and an adverbial ability to V by answering the question how to V?" (91).

Here is one way to see this. Habgood-Coote has offered a sufficient condition for the ability to V. There may be other ways to possess this ability,

but this one involves the ability to V in a specific, knowledge-involving way. Habgood-Coote offers this as an account of knowledge how, but how it fares in the conceptual space of the knowledge how debates is not my concern. For this also sounds (to me) like a covert account of knowledgeable action.

I said knowledgeable action is intentional action guided by knowledge of A-ing, and knowledge of how to A. The crucial part of my account is the explanation of how knowledge gets involved in the action, and of why it is specifically knowledge that is involved. I think this is just what Habgood-Coote's account needs, were we to think of it as an account of knowledgeable action. Early in his paper the question of knowledge arises. Habgood-Coote notices that one could find the right answer to a question of how to V deviantly (by guessing, for example). He proposes to put the issue to the side, instead singling out "the kind of answering that involves getting to the correct answer in the epistemically right way" (89). What way is this? Habgood-Coote writes "I will gloss this as the ability to know answers, since knowing an answer entails having got to the correct answer in the epistemically right (non-Gettiered, non-lucky) way" (89).

I'm aware that a man with a hammer sees too many nails. But again we see the importance of an account of non-deviant causation. In explicating knowledgeable action, such an account is crucial in being able to explain the controlled practical reasoning that provides entitlement to knowledge of A-ing and of how to A. Without it, we are forced to assume knowledge is involved, without an explanation. But an explanation is important, since agents are wily things and may very well find shortcuts and tricks enabling skill in the absence of knowledge. If it is knowledge specifically that we are interested in, it helps to understand how this knowledge comes about, and thereby helps to constitute this particular mode of agentive excellence.

8.10 Conclusion: Knowledgeable Action and Agentive Excellence

I have articulated a view on which the engagement of controlled practical reasoning in action embeds knowledge-constituting judgments regarding what the agent is doing, and how. The knowledge of action plays an important role in a process of rational control over that very action. I call action that involves knowledge of action in this way knowledgeable action. Knowledgeable action is a mode of agentive excellence. For in acting knowledgeably, the agent extends the reach of her control by way of her

controlled use of practical reasoning. She finds and implements good ways to execute her intentions (or desires or whatever), or to satisfy nearby background goals. The ways she finds, when she engages in controlled practical reasoning, are ways she might miss if she depended upon rigid routines, or upon whatever practice in the relevant domain might instill in her. This use of practical reasoning permits the exercise of control even in difficult circumstances—circumstances that render the agent uncertain, leading to practical reasoning.

There is a connection here with Anscombe's well-known claim that knowledge of action is in some sense practical as opposed to speculative. Here is Anscombe's central example of the practicality of an agent's knowledge:

> Imagine someone directing a project, like the erection of a building which he cannot see and does not get reports on, purely by giving orders. His imagination (evidently a superhuman one) takes the place of perception that would ordinarily be employed by the director of such a project. He is not like a man merely considering speculatively how a thing might be done; such a man can leave many points unsettled, but this man must settle everything in a right order. His knowledge of what is done is practical knowledge. (Anscombe 2000: 82)

For Anscombe, it seems that an agent's knowledge is practical at least in part because the agent relies on this knowledge to give his orders and to settle everything in the right order. So the agent knowledgeably determines the details of his project's execution.[15]

What I have done in this chapter is offer a fuller account of the way that knowledge is involved in the agent's execution of action, and of why it is knowledge that is involved, as opposed to something else. Knowledgeable action is a subset of intentional action: Anscombe over-generalized from the knowledge-involving cases. It is in knowledgeable action that the agent

[15] I am not here attempting to make full sense of the things Anscombe says regarding knowledge's practicality. Many have found Anscombe's view on these matters puzzling. Some representative remarks: "Despite its centrality to Anscombe's project, her concept of practical knowledge has not been widely understood" (Schwenkler 2015: 1). "Anscombe's account of 'practical knowledge' is hard to interpret" (Candlish and Damnjanovic 2013: 700) [Anscombe's view of practical knowledge is] "not just causally perverse but epistemically mysterious" (Velleman 2007: 103). "Anscombe is right to claim that there is something special about knowledge of what one is doing intentionally; but her explanation is not particularly helpful" (Setiya 2008: 392).

is well-placed to answer why-questions about what is going on. For in these cases the agents are posing those same questions to themselves, and answering them.

What distinguishes the actions that involve knowledge from those that need not? I suggest that the distinction can be drawn in terms of a requirement of practical reasoning for the execution of a certain (large) class of actions (and this requirement in turn seems to track a requirement upon flexibility in action execution). This requirement may stem from the complexity inherent in an action, or from difficulties inherent in the environment within which an agent acts. Either source may generate difficulties that agents need to work through as they carry out an action plan. When agents need to work through difficulties, they begin to reason practically, and at this point knowledge is generated and engaged.

This is an empirical claim. I and others have defended it (Shepherd 2015b; Christensen et al. 2016). Agents often need to engage practical reasoning at some point as they execute plans. This is a function of the complexity of human plans, the human behavioral control apparatus, and the nature of the circumstances in which humans operate. The result is that knowledgeable action, while a subset of intentional action, is nonetheless pervasive.

This is a different explanation for the pervasiveness of knowledge of action than one offered by Velleman (2007). Velleman maintains that agents have strong and persistent ("intrinsic") desires to know what they are doing and why they are doing it, and that (roughly) these desires motivate them to deliberate and act in ways that generate this knowledge. My explanation is simpler. Whether agents have such desires or not—and whether these desires play a crucial role in deliberation and action or not—I claim that knowledge of action is pervasive because execution of action usually requires the practical reasoning that generates it.