

Practical Knowledge First

Carlotta Pavese, forthcoming in *Synthese*

August 4, 2022

1 Introduction

Let us start with three examples.

Free Form According to a credited reconstruction of Pollock's first 'drip' painting, *Free Form* (1946), Pollock began by painting an entire canvas red and then added black and white tangles and pools of paint by flinging and dripping diluted oil paint from a brush or stick. Suppose my 5 year-old niece, armed with a canvas, a brush, and white, black, and red oil colors, started dripping color on the canvas randomly and without following a plan. While it is, presumably, not impossible for my niece to drip oil on a canvas in such a way that the outcome turns out to look exactly like Pollock's *Free Form*, if she did manage to reproduce, by pure luck, the exact same colors and forms on a canvas, the difference between her so doing and Pollock's performance would be no less noteworthy. Indeed, even if their bodily movements turn out to be exactly alike, Pollock's execution would be intentional throughout; my niece's would instead be accidental, fortuitous, and unintentional. For it seems clear that while my niece might intentionally move her brush in certain ways when dripping colors all over the canvas, she does not intentionally reproduce the exact same forms and colors on the canvas.

Fontana Lucio Fontana's highly personal style involved covering canvases with layers of thick oil paint applied by hand and brush and by using a scalpel or knife to create great fissures in their surface. Here again while Bruce the cat might scratch a canvas, and in so doing they might accidentally reproduce Fontana's scratches, they are unlikely to intentionally reproduce scratches of the exact length and shape.

Klutz and clown Consider a clown, who takes a tumble for an audience and compare it with a klutz, who trips and tumbles inadvertently (cf. Ryle 1949, p. 33). Intuitively the clown tumbles intentionally, while the klutz does not.

In these three cases, two subjects perform the same action but only one of them does it intentionally. And in all three cases, the contrast seems due, at least in part, to whether the subject possesses knowledge of what they are doing: Pollock knew what he was doing when painting *Free Form*; by contrast, my niece would not know what she was doing when producing a replica of *Free Form*. The clown but not the klutz knows what they are doing when they tumble. Fontana but not Bruce the cat knows what he is doing when producing identical scratches.

This idea that what is distinctive of intentional performances is one's practical knowledge in it — i.e., knowledge of what one is doing while doing it — famously traces back to Anscombe ([1963] 2000). While many philosophers have theorized about Anscombe's notion of practical knowledge (e.g., O'Brien (2007), Setiya (2008), Thompson et al. (2011), Schwenkler (2019)), there is a wide disagreement about how to understand it. This paper investigates how best to understand practical knowledge and its content for it to play the desired explanatory role in a reductive theory of *intentional action*, of *intention-in-action*, and of *control-in-action*.

Although I will often refer to how Anscombe seems to have conceived of practical knowledge, my primary goal here is not at all exegetical. Certain central features of Anscombe's conception of practical knowledge (e.g., its non-observational and non-inferential character) will not figure centrally in my discussion, though I believe what I will say is compatible with them. Indeed, in two main ways, my project is antithetical to Anscombe's own project, at least as usually interpreted. For one thing, Anscombe seems to have been quite generally skeptical of *reductionism* in philosophy (e.g., Carroll (2009), Teh (2013), Singh (2020)) and in particular of the possibility of understanding intentional action fully in terms of practical knowledge.¹ By contrast, my goal here is to make progress on how practical knowledge is to be understood for it to play a role in a reductive theory of skillful and intentional action and other action-theoretical notions, whether or not the resulting notion aligns even approximately with Anscombe's own views and even if the project is at odds with Anscombe's anti-reductionism. Moreover, Anscombe embraced a form of *anti-mentalism*, which explains her lack of interest for pursuing questions about the nature of mental representations underpinning practical knowledge; by contrast, my discussion will presuppose that mentalism is a legitimate approach in psychology and the philosophy of mind, and so that it makes sense to ask in virtue of which mental representations one can count as practically knowing.

The plan for this paper is incremental. §2 makes some preliminary remarks on how practical knowledge is to be understood in order for it to figure in a reductive theory of intentional action. §3 revises the initial proposal to provide a reductive theory of intention-in-action. §4 argues that in order to understand control-in-action in terms of practical knowledge, we ought to appeal to structured *practical modes of presentation* of the sort I have outlined in Pavese (2015). §5 discusses the case of basic actions and putative counterexamples to practical knowledge first. §6 closes by drawing more general conclusions about the relation between knowledge and action.

2 Reductive versus Non-Reductive practical knowledge

Let **Practical Knowledge** be the general claim that *knowing what one is doing while doing it* is necessary for intentional action:

Practical Knowledge One intentionally ϕ s only if one knows what one is doing when ϕ -ing.

¹Some textual evidence for Anscombe's anti-reductionism can be found in Anscombe 1971, p. 137.

The main task in this paper is to understand what kind of knowledge *knowing what one is doing* amounts to — i.e., how best to characterize its content and its epistemic profile.

Before delving into these questions, let me mention one influential objection to Practical Knowledge that might seem to undermine the entire project. Several authors have argued against the plausibility of Practical Knowledge on the grounds that certain actions can be intentional even though one is not seemingly aware of what one is doing (Paul 2009, p. 5).² For example, consider somebody driving home on autopilot; or a cabinet-maker immersed in his craft; both the driver and the cabinet-maker might respond to what is to be done without forming beliefs about what they are doing (cf. also Vekony et al. (2020)). The objection can be addressed by weakening Practical Knowledge as requiring only that one is *in position to know what one is doing when ϕ -ing*:

Practical Knowledge-Weak One intentionally ϕ s only if one is in position to know what one is doing when ϕ -ing.

Following a standard construal of *being in position to know*, we can understand *being in position to know p* , for some proposition p , as amounting to it being true that if one were to believe p , one would know it. So cashed out, Practical Knowledge-Weak also understands intentional action in terms of practical knowledge, since it understands it in terms of *being position to know what one is doing*. The stronger Practical Knowledge is most plausible for intentional actions which are also *skillful*, as the examples with which I have started. In these examples, it is hard to deny the psychological importance of practical knowledge. But a weakening of this practical knowledge condition along the lines described might be necessary to deal with this variety of objections. Having noticed it, I will bracket the issue in what follows.

How are we to unpack the content of knowing what one is doing to provide a plausible condition on intentional action? Compare the following necessary conditions:

Practical Knowledge-Non-Reductive One intentionally ϕ s only if one knows that one is ϕ -ing intentionally.

Practical Knowledge-Reductive One intentionally ϕ s only if one knows that one is ϕ -ing.

Both Non-Reductive and Reductive unpack *knowing what one is doing* in terms of propositional knowledge — knowledge *that something is the case*. Non-Reductive is a common way of formulating a practical knowledge condition on intentional action (e.g., Newstead (2006), Rödl et al. (2007), Schwenkler 2019, pp. 93; Piñeros Glasscock 2019); and yet by comparison, Reductive is more appealing to any philosopher who aims to provide a non-circular account of intentional action. For if we try to analyze intentional action in terms of knowledge that one is acting intentionally, then we will have presupposed the notion of intentional action in the analysans.

For those who are generally skeptical of the possibility of non-reductive theories of intentional action, here are further reasons to prefer Reductive to Non-Reductive, which are independent of one's reductionist ambitions. The first reason has been recently highlighted by an exchange between Piñeros Glasscock 2019 and Beddor & Pavese (2021). As Piñeros Glasscock (2019) has

²This is by no means the only objection to a practical knowledge condition on intentional action. For example, Setiya (2008) argue that in case of basic actions, one might intentional perform an action without knowledge or belief that one is doing so. I bracket these objections here. See Pavese (2021a) and Pavese (2020b) for discussion of these issues.

pointed out, it can be shown that, given plausible assumptions about knowledge (in particular, given the assumption that knowledge obeys a margin for error principle), Non-Reductive gives rise to absurdity, through an argument that resembles Williamson (1997)'s anti-luminosity argument.³ Piñeros Glasscock 2019, pp. 1248-9 starts by offering the following case:

Cleaning Sisyphus. The floors of the Underworld are getting filthy with blood and bile. Hades notices, and decides to give Sisyphus a more useful task than pushing a rock, repeatedly, up a hill. Handing him a mop, he orders him to clean the floors using the waters from the river Acheron, the cleanest in the Underworld. There is one problem: while at noon the waters of Acheron are clean as a spring, they slowly and gradually get dirtier and dirtier as the day goes by — by midnight it is just filth, much dirtier than the floors of the Underworld. Hades thus tells Sisyphus that at midnight he will be punished proportionally to his efficiency: his punishment will be worse if he fails to keep mopping when the water is still clean enough or if he keeps mopping when the water is dirtier than the floors. Each day, therefore, Sisyphus grabs his mop in the morning trying to mop as much of the Underworld as he can. . . . At noon, he is fully confident that he is intentionally acting under the description cleaning the floors. As the day goes by, however, he loses more and more confidence. But by midnight in those days when he keeps mopping to defy Hades, he is certain that he is not cleaning the floors of the Underworld (but instead making them dirtier).

From the description of the case, we have:

- (1) Sisyphus is intentionally cleaning the floors at t_0 (noon).

From (1) and Non-Reductive, we get:

- (2) Sisyphus knows he is intentionally cleaning the floors at t_0 .

Assume:

MARGIN FOR ERROR: If S knows that S is in condition C in a case σ , then in all nearby cases σ , S is in C.

A special instance of it together with (2) yields:

- (3) Sisyphus is intentionally cleaning the floors at t_1 .

By another application of Non-Reductive we get:

- (4) Sisyphus knows that he is intentionally cleaning the floors at t_1 .

Another application of Margin for Error lets us conclude that Sisyphus is intentionally cleaning the floors at t_2 . By repeated applications of this mode of reasoning, we arrive at the conclusion that at midnight, Sisyphus is intentionally cleaning the floors of the Underworld. But this contradicts the

³For worries about the margin for error principle, see Berker (2008). For recent defenses of a margin for error principle for knowledge, see instead Srinivasan (2015).

stipulation that he is not cleaning the floor at midnight. So Non-Reductive combined with Margin for Error leads to a paradoxical conclusion.

As Beddor & Pavese (2021) note, however, Reductive can be shown to escape this fate. From (1) and Reductive we can no longer derive (2), but only:

(5) Sisyphus knows he is cleaning the floors at t_0 .

Next, we might appeal to the relevant instance of MARGINS FOR ERROR and from (5) we can derive:

(6) Sisyphus is cleaning the floors at t_1 .

But, crucially, we cannot draw the stronger conclusion that Sisyphus is intentionally cleaning the floors at t_1 . So we cannot use Reductive to derive (3) from (6). Without (3), the *reductio* does not get off the ground. As they argue, the same *reductio* can be revived with Reductive only if further assumptions are granted — i.e., that there are such things as *essentially intentional actions* — which, however, are independently controversial (cf. Beddor & Pavese (2021)). So, while Non-Reductive conjoined with MARGIN FOR ERROR yields an absurdity, Reductive conjoined with MARGIN FOR ERROR is paradox-free.

There are two additional reasons to prefer Reductive to Non-Reductive. One source of trouble for Non-Reductive comes from non-human animals and small children. Presumably, these agents *can* act intentionally: Fido the dog might intentionally hunt for the bone that he buried. By Reductive, it follows that in the course of his search, Fido knows what he is doing — namely, searching for a bone. This much seems plausible. But does Fido know that he is intentionally searching for a bone, as Non-Reductive would have it? This seems doubtful. Presumably, in order to know that one is acting intentionally, one needs to have the concept of intentional action. But this seems to attribute an overly sophisticated theory of mind to our canine friend.

Finally, even if we restrict our attention to agents who possess the concept of intentional action, Non-Reductive faces counterexamples that do not arise for Reductive. Consider the following sort of case:

Neurological Meddling. Mary is a basketball player, and the stakes are high. So high that Sara, a devoted fan and rogue neurosurgeon, has implanted a control device in Mary's head. Sara will be closely watching the game. If she predicts that Mary will miss a shot, she will activate the control device and initiate a sequence that automatically controls Mary's movements. So as not to disorient Mary, Sara has carefully designed the device so that it gives Mary the illusion of control: even when it is activated, Mary will still experience the phenomenology of intentionally acting. Moreover, Mary knows all of this. The game commences: Mary takes a shot. As a matter of fact, the device was never activated.

It seems undeniable that in Neurological Meddling, Mary can intentionally take a shot. While doing so, she knows that she is taking a shot. But she does not know that she is taking a shot intentionally, because — for all she knows — the device might be activated. In this case, only Reductive delivers the right prediction.

3 Intention-in-action as practical knowledge

So there are a number of reasons for thinking that Non-Reductive is too strong a condition on intentional action. But while Reductive offers a more plausible condition, it can only provide *at best* a necessary condition — not a both necessary and sufficient one. In order to see why, consider again our earlier examples. If, e.g. in **Free Form**, my 5 year-old niece were shown **Free Forms** while she is busy covering the oil canvas with random colors, she might thereby realize that she is painting something very similar to that painting. However, this by itself does not make her reproduction of that masterpiece intentional. Or consider **Klutz and Clown**. It is true that while tumbling, the clown knows that they are tumbling and not so with the klutz. However, consider a case in which mid-tumble, the klutz realizes with consternation that they are tumbling. So both the clown and the klutz know that they are tumbling but it still seems true that only the former does so intentionally.

Ryle 1949, p. 33 originally used **Klutz and Clown** to illustrate a point that several action theorists later made — i.e., that intentional action requires skill or know-how (Bratman (1984), Mele & Moser (1994)). According to this idea, the difference between the clown and the klutz is that the clown arrives at their knowledge that they tumble in virtue of exercising their knowledge of how to tumble. By contrast, the klutz does not arrive at this knowledge in virtue of exercising their knowledge of how to tumble. Similarly, for my niece and Pollock. This observation motivates:

Epistemic theory One intentionally ϕ s just in case one knows that one is ϕ -ing in virtue of knowing how to ϕ .

However, Epistemic Theory will not also work as it stands, for it fails to capture one key aspect of our cases. It is not just the know-how that distinguishes, say, Pollock from my 5 year-old niece, or Ferrano from Bruce the cat, or even the clown from the klutz. To see this, consider the following variation on Ryle's case:

The Unlucky Clown As the clown enters the stage, just before they start their performance, they ought to walk down a long staircase. In their hurry to get to the stage, they stumble over their long clown pants and fall down. As they have intentionally produced this movement before, and thanks to their attentiveness to their movements they developed through their skill at being a clown, the clown does not take much to realize that they are tumbling down the stairs, right onto the stage in front of the public.

In this example, the clown comes to know that they are tumbling down the stairs in virtue of their skills as a clown (their having done it before and their clown developed sensibility allowing them to recognize their bodily movements as a case of tumbling). And yet intuitively their tumbling is not anymore intentional than in the case of the klutz. What the unlucky clown is lacking is *an intention*. The lack of intention is common to the other examples too. While Pollock draws his painting *by intending to do so*, my 5 years-old niece does not have the same intention. Similarly, while the clown tumbles *intending to tumble*, the klutz does not tumble with an intention to tumble. My niece and the klutz lack, we might say, an *intention-in-action*. Intentions-in-action are not quite the same as distal (or future-directed) intentions. A distal intention is an intention to perform an action in the future, not necessarily here and now (Bratman (1984)). By contrast, intentions-in-action are intentions that govern a performance *while it is happening*. So while a distal intention is not factive

(in the sense that one might distally intend to ϕ at some time t without actually ϕ -ing at t), if one intends-in-action to ϕ at t , one must be ϕ -ing at t . Indeed, plausibly, if one intends-in-action to ϕ , one also intentionally ϕ s.

This intention-in-action component is currently missing from Epistemic Theory. On the other hand, it is widely agreed that *some* intention-in-action is necessary for intentional action. If this is correct, we will not be able to provide a reductive theory of intentional action without also invoking intention-in-action. We could add a further condition to Epistemic Theory, requiring that the subject intends to perform the action. But the resulting theory would be more elegant if this intention condition were to be captured by an appropriate understanding of practical knowledge.⁴ So, let us explore the prospect of reducing intention-in-action to practical knowledge.

This is a juncture at which looking at what Anscombe originally said about practical knowledge might help us make some progress. Anscombe [1963] 2000, p. 57 famously claims that we cannot get a grip of practical knowledge without getting a grip of *practical reasoning* first. As several commentators of this passage have observed, Anscombe seems to have thought that practical knowledge and practical reasoning are closely connected, in that the former is based on, or constituted by, practical reasoning. From this remark of Anscombe's, some have even drawn the conclusion that practical knowledge is not propositional at all. After all, a piece of reasoning is not the same as a proposition; nor is the conclusion of practical reasoning (an intention or an action) obviously propositional either (cf. Frost 2019, p. 322). However, the view that practical knowledge is not a propositional state is at odds with Anscombe's own characterization of practical knowledge at several points, where she consistently uses that-clauses to characterize the content of practical knowledge. (Anscombe [1963] 2000, section 28, 50:2; 51:2; section 29, 51:3). In the light of this textual evidence, an interpretation of Anscombe's remark that is consistent with practical knowledge being propositional is, on the whole, preferable.

Here is a suggestion. Anscombe [1963] 2000, p. 80 says practical reasoning 'describes an order which is there whenever actions are done with intentions'. The suggestion that we cannot get a grip on practical knowledge if not by first getting a grip of practical reasoning might be understood as the claim that the content of practical knowledge somehow reflects the consideration of means to ends that is distinctive of practical reasoning. Thus, on this view, practical knowledge ought to be understood in such a way as *to result from the considerations of reasons for actions that is at the center of practical reasoning*. In the light of this, intentionally acting does not merely require knowledge *that one is acting*, as Reductive states. It requires knowledge that one is doing something *as a means to reach some further end*. This suggests that we should cash out Practical Knowledge as follows:

Reductive* One intentionally ϕ s only if when ϕ -ing, for some action ψ , one knows that one is ψ -ing in order to ϕ .

It seems clear that one cannot possess this knowledge without possessing an intention-in-action. For, arguably, one cannot perform some action in order to ϕ without possessing an intention-in-action to ϕ .

⁴Anscombe often emphasizes that practical knowledge causes the success of the action and in this respect shares intentions' direction of fit. By thinking of practical knowledge as closely related to intention-in-action, we might be able to account for the idea that practical knowledge is the cause of the action that it represents, and so accounts for its having *both* directions of fit.

If we cash out the content of practical knowledge in these means-to-ends terms, we can overcome the sort of counterexamples to the Epistemic Theory above. The unlucky clown does not know that they are performing certain bodily movements in order to tumble, since the clown is *not* performing those movements *in order to tumble*. My niece lacks knowledge that she is performing her color dripping in order to paint the canvas thus since it is false she is dripping the color on the canvas in order to paint it thus; *mutatis mutandis*, for Bruce the cat in the Fontana’s example. They all lack both intention-in-action and practical knowledge.

The current suggestion is that an intention-in-action is entailed by one’s knowledge that one is doing something as means to reach some further ends. Now, at this juncture, there are two theoretical options ahead of us. On one option, intention-in-action *is* the object of practical knowledge, in the sense that having practical knowledge of what one is doing is a matter of knowing that one is intending-in-action. The problem with this suggestion is that the sort of anti-luminosity worries raised by Piñeros Glasscock (2019) might resurface here, since this view would basically amount to saying that intention-in-action is a luminous condition, such that whenever one is intending-in-action to ϕ , one would practically know that one does.⁵

The alternative theoretical option which I want to explore takes intention-in-action *to consist* in practical knowledge in the sense that for one to intend-in-action to ϕ is for one to know, when ϕ -ing, that one is ψ -ing in order to ϕ , for some means ψ :

Epistemic theory* One intentionally ϕ s just in case when ϕ -ing, one knows that one is ψ -ing in order to ϕ for some action ψ , in virtue of knowing how to ϕ .

According to Epistemic theory*, intentional action reduces to intention-in-action, through practical knowledge. This view does not face any of the problems that Non-Reductive faces. It is true that, in a sense, Epistemic Theory* *embeds* the structure of practical reasoning into the content of the practical knowledge, by taking this content to be structured in terms of means and ends. However, the over-intellectualization objection that afflicts Non-Reductive does not apply to it. Although animals and children most plausibly lack a fully developed theory of mind that is presumably required for possessing the concept of intentional action, they might well have a rudimentary conception of means and ends. Secondly, Epistemic Theory* can account for the judgment that in Neurological Meddling, Mary’s shot is intentional, since although she does not know that she is taking the shot intentionally, Mary presumably knows that she is performing this and that bodily movement in order to take a shot at the basket. Finally, Epistemic Theory* does not trigger an anti-luminous regress, since according to it, practical knowledge is not knowledge of any mental condition; so Epistemic Theory* does not amount to claiming that any mental condition is luminous. So just as with Reductive, the vicious regress can not take off the ground.

4 Practical knowledge, control, and practical senses

§2 started with a naive account of knowing what one is doing in terms of Reductive, which was then refined in §3 in order to provide an account of both intentional action and intention-in-action.

⁵I think the converse (that if one intentionally ϕ then one intends-in-action to ϕ) might not hold true, due to the sort of counterexamples to the so-called Simple View (i.e., that intentionally ϕ requires intending to ϕ) which Bratman et al. (1987) raised. I cannot discuss this complex issue here, so I will bracket it.

In this section, I suggest a further refinement in order for it to account for *control-in-action* — a property theorized by many cognitive scientists as well as philosophers of psychology as the distinctive property of *skillful behavior*.⁶

According to many scholars, control-in-action requires something like a monitoring condition on one's action. Different philosophers characterize this monitoring condition differently. So, for example, some have argued that one is in control of one's action if one's attention and intention to perform it couple in the right way, throughout the whole performance (Wu (2016)). However, it is doubtful that every controlled action is attended (see Pavese (2021c) for more discussion). Consider, e.g., intentional omissions: one might intentionally omit doing the dishes all afternoon, but that does not involve attending to one's laziness through the afternoon. So while many controlled actions are attended, attention is unlikely to be the sort of general monitoring condition that we are looking for when characterizing control. By contrast, Practical Knowledge seems to provide a more general monitoring condition on intentional action, as it does not need to involve conscious attention. So even in intentional omissions — when e.g., I am intentionally not washing the dishes throughout the afternoon — one might know, or be in a position to know, what one is doing throughout the afternoon. The question I would like to discuss in this section is how practical knowledge is to be understood in order to capture control.

As characterized in §3, knowing what one is doing consists in *knowing that one is performing some action in order to achieve some further goals*. The proposal is that one intentionally acts if one knows that *throughout* the performance of the relevant action. This way of characterizing practical knowledge is too *static* for it to enter into a satisfactory theory of control-in-action. For being in control during a performance presumably requires *updating one's cognitive state while performing the action*. Moreover, as I am performing a task, I might know that I am doing it in order to achieve some further ends; but that does not make me fully in control of what I am doing. For example, I might know that I am answering a questionnaire in order to take an exam. However, I am not necessarily in control of taking an exam (for example, if I am unprepared). Intuitively, control-in-action requires one knowing, *for each part of the task why one is performing it at the moment one is doing it*. The current understanding of practical knowledge, as outlined in §3, is insufficiently dynamic to capture control-in-action.

Here is a suggestion. Being fully in control of one's action is not simply a matter of one being in a static practical knowledge state — as the Epistemic Theory* has it. Something stronger has to be true. It has to be the case that for every part of the task, one knows that one is performing that part as a means to make progress in the completion of the task. As the task unfolds, one executes different parts of that task. So as the task unfolds, one would have to know, or being in position to know, for different parts, that those are means for the next part of the task. That is to say that in order to understand the role of practical knowledge for control-in-action, we ought to understand

⁶Christensen et al. (2016) argue that cognitive control is present in expert skillful performance in the form of 'higher-order action control'. Wu (2016) focuses on *agentive control*, which he thinks 'yields phenomena of central philosophical interest: moral, rational, reason-based, skilled, conscious, epistemic and free agency' (Wu, 2016, p. 101). Verbruggen & McLaren (2014) argue that adaptive and goal-directed behavior is to be understood in terms of agentive or executive control, in virtue of which people can regulate their behavior according to higher-order goals or plans. Finally, Toner et al. (2015) argue that skilled athletes use cognitive (or executive) control to maintain or improve their performance proficiency (see also Christensen et al. (2015), Fridland (2014), Christensen et al. (2019), Shepherd (2014), Wu (2011), Shepherd (2021)). These psychologists and philosophers all seem to consider it to be a hallmark of skillful performance that it is, in some sense to be specified, under the agent's control.

it as a *dynamic* state that guides the execution of the action from start to finish.

This dynamic construal of *knowing what one is doing* can be accounted for by invoking *practical modes of presentation* as integral to the know-how component of the Epistemic Theory*. In the following, I will first motivate the appeal to practical modes of presentation in a theory of know-how; then, I will put them to use in a theory of control-in-action.

I can wiggle my left ear back and forth. It is a skill that I have. I have done it several times, and so I can intend to do it again. Indeed, arguably I can come to know that I will wiggle it by deliberating to wiggle my left ear back and forth. But when I choose to wiggle my ear back and forth, I do not come to know that I will perform that act *under a physiological mode of presentation* — I might have no idea what the relevant muscles are and what I am doing with them. Nor do I need to come to know that I will perform that act *as the act of wiggling one's ear* (I might not have learned the English expression 'wiggle one's ear' yet, or any corresponding expression in any language). *Demonstrative ways of thinking* will not cut it either. For in this example, I can think of my wiggling one's ear without performing the action, so without the referent of a relevant demonstrative way of thinking being brought into existence. Finally, I do not know that I will perform that act *under an observational concept either*, for I might be able to think of that act even in absence of perceptual feedback, or even in absence of proprioception. And yet if I can know that I will wiggle my ears in any way, I must be able to identify and think of my wiggling of my ear in some way — in a special practical way.

Thinking practically in this fashion is to represent through practical modes of presentation. The appeal to practical modes of presentation is due to intellectualists about knowing-how, who embrace the view that knowing-how is a kind of propositional knowledge. The claim that know-how is a kind of knowledge-that encounters an immediate incredulous stare. One outstanding objection is the *sufficiency objection*: how could propositional knowledge be sufficient for knowing how to do something? (cf. Ryle 1949, p. 5) How could, the sufficiency objection goes, one know how to perform a task just in virtue of knowing a proposition about how to perform it? Here is an obvious counterexample:

Swimming: Suppose I look at a swimmer's swimming, and my swimming instructor, pointing to the swimmer, says to me, "That is a way in which you could swim too." I believe my instructor and we may suppose that what she said is in fact true. I may thereby come to know a true answer to the question "How could I swim?" However, in the relevant sense, I may not have come to know how to swim. If I took a swimming test, I might still fail it. If thrown in the swimming pool, I might still drown. I do not know how to swim in the relevant sense and yet I do know a true proposition about how to swim.

In response to this sort of counterexample, intellectualists argue that knowing a proposition observationally or demonstratively is not the same as knowing it practically (Stanley & Williamson (2001), Pavese (2015), Pavese (2019), Pavese (2020a)). For one to know how to swim, in the relevant sense, one must know of a way to swim represented under a distinctive practical mode of presentation, which is essentially different from the observational or demonstrative mode of presentation in **Swimming**. This kind of practically represented propositional knowledge is what (some) intellectualists call know-how and is absent in the example above.

Some intellectualists have explored ways of responding to the sufficiency objection that do not appeal to practical modes of presentation. For example, Stanley (2011) argues that practical modes

of presentation can be dispensed with by intellectualism about know-how and considers answering the sufficiency challenge in cases such as **Swimming** by appealing to the context-sensitivity of the ability modal “could”. According to Stanley, depending on how the context for the modal is restricted, “That is how you could swim” could mean either that that is how you can swim given your current physical state or that that is how you could swim after training. But coming to know that that is how I could swim after training is clearly not enough for me to come to know how to swim now. Instead, the argument goes, what one needs to know is the former proposition: that that is a way to swim given my current physical state. Other intellectualists have endorsed this response (cf. Cath (2019), Cath (2020)). So they conclude that practical modes of presentation are not needed for intellectualism about knowing-how.

Yet, there is reason to doubt that intellectualism can overcome the sufficiency challenge by dispensing altogether with practical modes of presentation. Knowledge of abilities is not sufficient for know-how either — however restricted the scope of the ability modal is. Consider a variant of the previous scenario but where Mary is a skilled swimmer who is, one day, affected by memory loss and so forgets how she is able to swim. Nothing has changed in Mary’s physical state: she is still able to swim but she has just forgotten how she is able to swim. Suppose she is told, by looking at a recording of her swimming the day before, that that is how she can in fact swim given her current physical state. She might come to know how she is in fact able to swim (just like that!). Yet, she would still fail to know how to swim in the relevant sense and still drown if thrown into the pool. This is so because failing to be able to identify that way of swimming practically, she cannot act on it when needed. The right conclusion to draw is that knowledge of abilities — even suitably restricted — does not suffice for know-how.

Practical modes of presentation help capture another central feature of know-how. Know-how is not always *articulable*: skilled agents often cannot report how they perform the tasks they are skilled to perform (Schiffer (2002), Noë (2005)).⁷ Unarticulability is not surprising if knowledge-how involves practical modes of presentation. For practical modes of presentation differ from descriptive (or semantic) modes of presentation of the same tasks precisely in that they are not necessarily verbalizable (cf. Pavese (2021d)). For one might possess the practical mode of presentation but not the semantic mode of presentation of that task, and hence lack semantic articulable knowledge about the task.

Practical modes of presentation also enter fruitfully in a theory of practical knowledge. After all, a theory of intentional action and control in terms of practical knowledge *also* faces the *articulability* objection. One might act intentionally without being able to explain how one does so. A classic example of this phenomenon are chicken sexers, who can intentionally find out the sex of chickens and yet lack the ability to explain to others how they do it. If intentional action requires

⁷Stanley 2011, p. 161 responds to the articulability objection by pointing out that there is a sense in which propositional knowledge is always verbalizable. A punch-drunk boxer, who can demonstratively refer to his re-enactment of the way of boxing against southpaws, and says, “This is the way I fight against a southpaw,” intuitively knows that this is the way he fights against southpaws. This knowledge has an essential demonstrative component. But the same goes for much of other propositional knowledge, for example the knowledge we express by saying, “This is the tool for the job,” or “That is going to be trouble.” So, the response goes, know-how is verbalizable just like propositional knowledge is supposed to be verbalizable. However, this reply to the articulability objection assumes that ways to execute tasks are always ostensible and as such can always be picked up by a demonstrative. This does not need to be so: on any single occasion, one may only act on *parts of a way*. So for example, consider a swimmer who knows how to swim, and so, according to intellectualism, knows a proposition about how to swim. And yet when he swims, he only swims in a particular way, that adapts well to the circumstances in which he finds himself.

practical knowledge as understood in §3 — a sort of propositional knowledge of means to ends — how can this be? If when knowing what one is doing, one represents what one is doing practically, rather than semantically or observationally, then one might know what one is doing and yet lack the ability to explain what one is doing.

The next point is that, correctly construed, practical modes of presentation explain the importance of practical knowledge for control. For our purpose, it will be helpful to have a concrete theory of practical modes of presentation on the table. According to Pavese (2015), representing a task practically means *representing it in terms of the subject's most basic practical abilities*. This notion of practical representation can be defined inductively, first for basic actions and then for complex actions. If the task is a basic action, to practically represent a basic action is simply to represent it *in a way that enables the agent to perform that action when they intend to do so*; to practically represent a complex action, instead, is *to represent it through a procedure that is effective for that agent*. Here an effective procedure for an agent is a way of breaking down the task in terms of the most basic operations that the agent can execute and in terms of basic modes of combination of those operations.

Now suppose one represents a task practically in this sense, by knowing how to perform it. Suppose, moreover, one intends to perform this task and executes this intention in action. At each time during the execution of the task, one is in position to know that one is performing that particular part of the task as a means for executing the whole task. So, at each point in time, one is in position to know that one is performing a part of the task towards the goal of performing that task. And the task unfolds, one practically knows that one is performing the next part of the task, and so on. In virtue of practically representing a task — i.e., in virtue of representing it by breaking it down into parts that one can perform — one is in position to know at each point of the execution that one is performing some or other part of the task and so, one is continuously in control of the execution of the task.

This accounts for the stronger notion of control that we started with. For control, the content of practical knowledge has to be understood as representing each part of the task in a way that one can execute it — i.e., one must represent each part of the task practically. So as I am taking an exam, I am in control of my so doing if I know, as the exam unfolds, that I am performing that part of that exam thusly in order to execute the next part. This dynamic knowledge state does seem to be the sort of cognitive state that can account for control in action.

In conclusion, by combining practical knowledge and practical modes of presentation understood as part-whole structured practical representations, one can provide a satisfying model of practical knowledge and of the control that it affords, as well as a response to the articulability challenge. Appealing to practical modes of presentation in order to cast light on control comes with further theoretical benefits. According to Epistemic Theory*, the practical knowledge the agent has of her own action is partially grounded on the know-how she possesses. The suggestion is that it is the exercise of a particular know-how that puts her in a position to know that she is doing such and such. But one might ask: how does the exercise of a particular know-how allow the agent to acquire that further piece of practical knowledge? Epistemic theory*, combined with the notion of control that I am offering, allows to shed better light on this 'in virtue of' relation: practical knowledge is at least partly grounded on knowledge-how, since knowledge-how involves practical modes of presentation, which also figure centrally in the dynamic state in which practical knowledge consists.

5 Basic actions and chancy actions

A common objection to a practical knowledge condition on intentional action is that it is too strong. Consider the case of basic actions: if an action does not need any other action as its means, in what sense does it require practical knowledge? Or consider the case of intentional, though chancy, actions: is practical knowledge really needed in these cases?

Applied to basic actions, Epistemic Theory* would require that for one to perform a basic action ϕ intentionally, one would have to know, for some ψ , that one is ψ -ing in order to ϕ . But basic actions are actions that can be performed intentionally without thereby performing any other action. So what action ψ is such that one needs to know that one is ψ -ing in order to ϕ ? Different answers are available here but I will focus on one that follows naturally from the discussion of the previous section. We could maintain that every basic action is *its own means*. According to this view, basic actions only require trivial means. On the resulting view, someone intentionally lifts their finger provided their action is guided by their knowledge that one is lifting one's finger in order to lift one's finger. This raises the question: how could this trivial knowledge guide their action? The answer is that the agent needs to have a practical way of representing the action in question — that is, a way that enables the agent to perform the action when they intend to do so. Recall my wiggling my ear back and forth. It is a basic action that I can do on demand. But for me to be able to select it among many possible others, I need to be able to think of it in a way that is practically relevant. This holds for complex as well as for basic actions. For example, suppose a doctor asks me to do lift my finger in front of theirs. In order to comply with the instruction, I need to be able to represent my finger, from the inside, so to say, and I need to do so in a way that will enable me to lift it. This suggests an answer to the question of how seemingly trivial knowledge can guide the agent. The relevant knowledge is not simply the (trivial) knowledge that one is ϕ -ing in order to ϕ — rather it is the knowledge that one is [ψ -ing] in order to ϕ — where the bracket stands for a practical way of identifying the action by means of which a subject can perform the action when they intend to do so.

The second issue concerns intentional but very chancy actions. When actions are very lucky, it is unclear that for them to be intentional, one ought to know that one is performing them. Since, after all, it is very unlikely that the performance will be successful. Consider the following putative counterexample, proposed by Shepherd & Carter (2021):

Baseball

Here is a case involving baseball that brings the right issue to light. For those not versed in baseball, we focus on batting. The batter stands 60 feet, 6 inches from a mound upon which stands the pitcher. The pitcher throws a small white ball towards the batter, and the batter attempts to hit the ball with a bat. This is not always easy, since the pitch comes in at varying speeds (often anywhere from 80 to 100 mph), with different kinds of spin and curvature. Say that the greatest hitter of all time (call him Pujols) approaches the plate and forms an intention to hit a home run — that is, to hit the ball some 340 feet or more in the air, such that it flies out of the field of play. Say that Pujols also believes that he will hit a home run, and that he has the practical belief as he is swinging, that he is hitting a home run. As it happens, Pujols's behavior, from setting his stance and eyeing the pitcher, to locating the pitch, to swinging the bat and

making contact with the ball, is an exquisite exercise of control. Pujols hits a home run, and so his belief that he is doing just that is true.

Shepherd & Carter (2021) argue that Pujols intentionally hit a home run. And yet he does not know at any point during the action that he will hit a home run. They conclude that practical knowledge is not necessary for intentional action. But of course this is no objection to a practical knowledge condition on intentional action, when appropriately formulated as in Reductive* and Epistemic Theory*: while Pujols might not know that he is hitting a home run, he can definitely know that he is taking some means to that end — e.g., that he is swinging the bat thus and so in order to hit the ball, that he is hitting the ball in order for it to fly out the field of play, that he is touching bases in order to hit a home run, etc. So, Pujols can have practical knowledge in the sense required for Epistemic Theory*. And as we have seen in §3, the way of cashing out knowledge of what one is doing offered by Epistemic Theory* is independently motivated by the need to provide a reduction of intention-in-action.⁸

6 Conclusions: Knowledge-First and Action Theory

The so-called knowledge-first program has been primarily construed as a program in epistemology and the philosophy of mind. In epistemology, knowledge-firsters attempt to understand epistemological notions such as justification and evidence in terms of knowledge (e.g., Williamson (2000), Aarnio (2010); Littlejohn et al. (2017)); in philosophy of mind, they have emphasized the role of knowledge-representation in folk psychology (Nagel (2017), Phillips et al. (2021)). Yet another chapter of the knowledge-first program — one which has only been partially written — belongs squarely to the philosophy of action and looks at the role of knowledge in explaining intentional and skillful action (e.g. Gibbons (2001), Pavese (2021b), Pavese & Beddor (2022)).⁹ This essay has contributed a section to this new chapter of the knowledge-first program — a section that focuses narrowly on the role of practical knowledge in action theory. I have argued that practical knowledge can provide a reduction of skillful, intentional action, and intention-in-action to a dynamic propositional knowledge state, which is best characterized by an appeal to structured

⁸It is also worth noting that in the recent literature, others have noted independently (cf. in particular Setiya 2007, p. 25) that even Davidson's carbon copier objection, and similar other cases, can be satisfactorily addressed when practical knowledge is understood along these lines. In Davidson's example, someone is writing heavily on a page, intending to produce ten legible carbon copies (Davidson (1978)). While writing, the agent does not know that they are producing ten legible carbon copies. Still, if they succeed in their goal, it seems that they did so intentionally. According to Epistemic Theory*, that is so because although the carbon copier might not know that they are producing ten carbon copies, they must know that they are writing firmly on a stack of copies in order to produce ten carbon copies. So the practical knowledge condition is satisfied. Indeed, the standard way of dealing with Davidson's carbon copier objection (cf. in particular Setiya 2007, p. 25) already provides a satisfactory response to **Baseball**.

⁹An interesting question that I have not attempted to address here is whether an epistemic theory of intentional action is more defensible than the more familiar commitments of knowledge-first epistemology. An anonymous referee notes that some of the major challenges to Knowledge-First epistemology (as traditionally construed) do not obviously generalize to an epistemic theory of intentional action. For example, Brown (2008) has challenged knowledge norms of practical reasoning using high stakes cases; Roeber (2018) has argued that these challenges can be generalized beyond high stakes scenarios. While I am not myself fully persuaded by these challenges to the knowledge-first account of evidence, it may well turn out that an epistemic theory of intentional action is *more* defensible than the more familiar commitments of Knowledge First epistemology.

practical modes of presentation. Practical modes of presentation also enter centrally in an account of practical knowledge for basic actions. Finally, some recent counterexamples to a practical knowledge condition on intentional action misfire.

References

- Aarnio, M. L. (2010). Unreasonable knowledge. *Philosophical Perspectives*, 24, 1–21.
- Anscombe, G. E. M. ([1963] 2000). *Intention*. Boston: Harvard University Press.
- Anscombe, G. E. M. (1971). *Causality and determination: An inaugural lecture*. CUP Archive.
- Beddor, B., & Pavese, C. (2021). Practical knowledge without luminosity. *Mind*.
- Berker, S. (2008). Luminosity regained. *Philosophers' Imprint*, 8.
- Bratman, M. (1984). Two faces of intention. *The Philosophical Review*, 93(3), 375–405.
- Bratman, M., et al. (1987). *Intention, plans, and practical reason*, vol. 10. Harvard University Press Cambridge, MA.
- Brown, J. (2008). Subject-sensitive invariantism and the knowledge norm for practical reasoning. *Noûs*, 42(2), 167–189.
- Carroll, J. (2009). Anti-reductionism. *The Oxford handbook of causation*, (pp. 279–298).
- Cath, Y. (2019). Knowing how. *Analysis*, 79(3), 487–503.
- Cath, Y. (2020). Know how and skill: The puzzles of priority and equivalence. In E. Fridland, & C. Pavese (Eds.) *Routledge Handbook of Skill and Expertise*. New York: Routledge.
- Christensen, W., Bicknell, K., McIlwain, D., & Sutton, J. (2015). The sense of agency and its role in strategic control for expert mountain bikers. *Psychology of Consciousness: Theory, Research, and Practice*, 2(3), 340.
- Christensen, W., Sutton, J., & Bicknell, K. (2019). Memory systems and the control of skilled action. *Philosophical Psychology*, 32(5), 692–718.
- Christensen, W., Sutton, J., & McIlwain, D. J. (2016). Cognition in skilled action: Meshed control and the varieties of skill experience. *Mind & Language*, 31(1), 37–66.
- Davidson, D. (1978). Intending. In *Philosophy of History and Action*, (pp. 41–60). Springer.
- Fridland, E. (2014). They've lost control: Reflections on skill. *Synthese*, 191(12), 2729–2750.
- Frost, K. (2019). A metaphysics for practical knowledge. *Canadian Journal of Philosophy*, 49(3), 314–340.
- Gibbons, J. (2001). Knowledge in Action. *Philosophy and Phenomenological Research*, 62(3), 579–600.

- Littlejohn, C., Carter, A., Gordon, E., & Jarvis, B. (2017). How and why knowledge is first. *Knowledge First: Approaches in Epistemology and Mind*, (p. 19).
- Mele, A., & Moser, P. (1994). Intentional action. *Nous*, 28, 39–48.
- Nagel, J. (2017). Factive and nonfactive mental state attribution. *Mind & Language*, 32(5), 525–544.
- Newstead, A. (2006). Knowledge by intention? on the possibility of agent’s knowledge. In S. Hetherington (Ed.) *Aspects of Knowing*. Elsevier Science.
- Noë, A. (2005). Anti-Intellectualism. *Analysis*, 65(288), 278–90.
- O’Brien, L. (2007). *Self-knowing agents*. Oxford University Press.
- Paul, S. K. (2009). *How we know what we’re doing*. Ann Arbor, MI: Michigan Publishing, University of Michigan Library.
- Pavese, C. (2015). Practical senses. *Philosopher’s Imprint*, 15(29).
- Pavese, C. (2019). The psychological reality of practical representation. *Philosophical Psychology*, 32(5), 784–821.
- Pavese, C. (2020a). Practical representation. In *Routledge Handbook of Philosophy of Skills and Expertise*, (pp. 226–44). Routledge.
- Pavese, C. (2020b). Probabilistic knowledge in action. *Analysis*, 80(2), 342–356.
- Pavese, C. (2021a). Know-how, action, and luck. *Synthese*, 198(7), 1595–1617.
- Pavese, C. (2021b). Knowledge, action, and defeasibility. In J. Brown, & M. Simion (Eds.) *Reasons, Justification, and Defeaters*, (pp. 177–200). Oxford University Press.
- Pavese, C. (2021c). Knowledge and mentality. *Philosophical Perspectives (A Supplement to Nous)*, 35:1, 359–383.
- Pavese, C. (2021d). Practical concepts and productive reasoning. *Synthese*, 199, 7659–7688.
- Pavese, C., & Beddor, B. (2022). Skills as knowledge. *Australasian Journal of Philosophy*, forthcoming, 359–383.
- Phillips, J., Buckwalter, W., Cushman, F., Friedman, O., Martin, A., Turri, J., Santos, L., & Knobe, J. (2021). Knowledge before belief. *Behavioral and Brain Sciences*, 44.
- Piñeros Glasscock, J. S. (2019). Practical knowledge and luminosity. *Mind*, 129(516), 1237–1267.
- Rödl, S., Sebastian, R., et al. (2007). *Self-consciousness*. Harvard University Press.
- Roeber, B. (2018). The pragmatic encroachment debate. *Noûs*, 52(1), 171–195.
- Ryle, G. (1949). *The Concept of Mind*. Chicago: Chicago University Press.

- Schiffer, S. (2002). Amazing Knowledge. *Journal of Philosophy*, 99(4), 200–202.
- Schwenkler, J. (2019). *Anscombe's Intention: A Guide*. Oxford University Press.
- Setiya, K. (2007). *Reasons without rationalism*. Princeton University Press.
- Setiya, K. (2008). Practical knowledge. *Ethics*, 118(3), 388–409.
- Shepherd, J. (2014). The contours of control. *Philosophical Studies*, 170(3), 395–411.
- Shepherd, J. (2021). *The shape of agency: Control, action, skill, knowledge*. Oxford University Press, USA.
- Shepherd, J., & Carter, J. A. (2021). Knowledge, practical knowledge, and intentional action. *Ergo*.
- Singh, K. (2020). Anscombe on acting for reasons. In *The Routledge Handbook of Practical Reason*, (pp. 172–184). Routledge.
- Srinivasan, A. (2015). Are we luminous? *Philosophy and Phenomenological Research*, 90(2), 294–319.
- Stanley, J. (2011). *Know How*. Oxford: Oxford University Press.
- Stanley, J., & Williamson, T. (2001). Knowing How. *Journal of Philosophy*, 98(8), 411–444.
- Teh, N. J. (2013). Anscombe on non-reductionistic accounts of human action. *Willing the Good: Empirical Challenges to the Explanation of Human Behavior*, (p. 227).
- Thompson, M., Hornsby, J., & Stoutland, F. (2011). 7. Anscombe's intention and practical knowledge. In *Essays on Anscombe's intention*, (pp. 198–210). Harvard University Press.
- Toner, J., Montero, B. G., & Moran, A. (2015). Considering the role of cognitive control in expert performance. *Phenomenology and the Cognitive Sciences*, 14(4), 1127–1144.
- Vekony, R., Mele, A., & Rose, D. (2020). Intentional action without knowledge. *Synthese*, (pp. 1–13).
- Verbruggen, F., & McLaren, I. (2014). Banishing the control homunculi in studies of action control and behavior change. *Perspectives on Psychological Science*, 9(5), 497–524.
- Williamson, T. (1997). Knowledge as evidence. *Mind*, 106(424), 717–741.
- Williamson, T. (2000). *Knowledge and Its Limits*. Oxford, UK: Oxford University Press.
- Wu, W. (2011). Confronting many-many problems: Attention and agentive control. *Noûs*, 45(1), 50–76.
- Wu, W. (2016). Experts and deviants: The story of agentive control. *Philosophy and Phenomenological Research*, 93(1), 101–126.