

The Epistemology of Skills

Forthcoming in *Blackwell Companion to Epistemology*

Carlotta Pavese

1. Introduction

Artemisia Gentileschi was an exceptionally skilled painter. She undoubtedly had a knack for it, which she inherited from her father Orazio Gentileschi. But her talent alone did not make her into one of the most skilled Baroque painters—it took several years of training first under her father’s supervision, and then at the Tuscan Academy of the Arts of Drawing, where she was the first woman ever to enroll and where she learned from the best painters of her time. Consider her ability to paint at the culmination of her career, as is manifested in *Self-portrait as the Allegory of Painting* (1638-1639). In what sense does that ability count as a skill?

Skills such as that Gentileschi had are *actively learned* abilities that characteristically manifest in *controlled actions*. Skills are distinguished from bodily and cognitive instincts, from knacks, talents, as well as from general faculties in their being learned; the distinctive way in which they are learned distinguishes them from habits; their distinctive agentive control distinguishes skills from reflexes and incompetent behavior, whereas their voluntary control distinguishes skills from virtues and from epistemic competences. Skills are also closely related to know-how, as it seems that one cannot be skilled at a task without knowing how to perform it; on the other hand,

knowing how to perform a task might not suffice as a skill—after all, many unskilled painters might nonetheless know how to paint (just not as well as Artemisia Gentileschi).

If distinctive kinds of learning and control are central to paradigmatic cases of skills, then an epistemology of skills includes questions such as the following: what kind of learning does skill acquisition involve? How is learning of a skill even possible? What kind of knowledge, or know-how, if any, do skilled agents possess? How are we to think of this knowledge in such a way as to explain the control that is distinctive of skilled actions?

This chapter first motivates a novel demarcation of skills from other abilities on the basis of their robust and active learnability as well as of their distinctive kind of control, and then discusses some questions that arise for the epistemology of skills.

2. The Demarcation Question

Ordinary use of ‘skill’ is very liberal, and applies indiscriminately to knacks, talents, instincts, know-how, general faculties, as well as to innate or acquired abilities.

Professional philosophers’ use of ‘skill’ is often equally liberal. Many epistemologists classify epistemic competences as kinds of skills and ordinarily compare perception to archery and piano playing (e.g., Sosa and Callahan 2020) or to fencing (Stanley and Williamson 2017); some go as far as to think of knowledge as a skill of sorts (Heatherington 2020). Philosophers of mind and cognitive science talk of perception (Siegel 2020), of attention (Wu 2020), of memory (Goldwasser 2023), of imagination (Kind 2020) as of skills, and equate them to crafts such as carpentry or tool use.

These disparate uses of ‘skill’ are no doubt motivated by the relevant theoretical goals. But they have the pitfall of classifying different kinds of practical abilities under the same heading; by doing so, they risk obliterating significant differences. Some of these differences have epistemological import. My goal in this section is to highlight some important differences among different kinds of abilities and to single out a narrower, and technical, notion of skill that (i) has a long historical pedigree, (ii) happens to pick out a natural (i.e., psychological) kind, and (iii) it is also of special interest to epistemologists, as it is constitutively connected to learning and knowledge. As it turns out, skills stand out from general faculties such as perception, memory, reasoning, and attention as well as from instincts, virtues, powers, and habits in their characteristic manifestations and in the sort of control that their manifestations exhibit. They are also distinctive in their cultural dimensions and in the manner of their acquisition.

Some initial distinctions are uncontroversial. Paradigmatic skills such as archery, chess, or playing a musical instrument differ from physiological abilities, such as the ability to digest or to breathe; skills also differ from a variety of cognitive but subagential abilities, such as the ability to process relevant complex acoustic and/or visual inputs or the ability to store information in working memory. Skills differ from these abilities in that they characteristically manifest in *actions*—something agents do (cf., e.g., Ryle 1949; Pavese 2013). This is true of both theoretical skills (such as math, philosophy and natural sciences), intellectual skills (e.g. chess) and embodied skills (e.g., archery, basketball, gymnastics): embodied skills characteristically manifest in bodily actions—actions that require the use of muscles for their performance (in particular that

of striatic and skeleton muscles, or ‘controllable muscles’); intellectual and theoretical skills characteristically manifest in *mental actions*, which do not require the use of muscles for their performance (cf. Schmidt and Wrisberg 2008; Pavese 2024). By contrast, the characteristic manifestations of the ability to digest (i.e., digestion) and of the ability to breathe (i.e., breathing) are not purposeful and goal-directed, nor do they need to be under the full control of the agent. Saying that skills characteristically manifest in actions does not preclude them from manifesting *non-characteristically* in involuntary movements and in reflexes too. For example, a skilled basketball player might instinctively catch the ball even when doing so will thwart their aims. The catching still manifests their skill—it is definitely evidence that the skill is there—albeit not characteristically.

Skills allow for different stages of acquisition—novice, beginner, competent, proficient, and expert (Dreyfus S. E. 2004). Control is a fundamental feature of the expert level of performance. A skilled painter or a skilled pianist is in control of their execution when they exercise their skills. Their control is manifested in their ability to adjust the execution of the action as the performance unfolds. Agentive control—the sort of control characteristic of skilled behavior—has been theorized by cognitive scientists, as well as philosophers of psychology (Miller 2000, 59; Dreyfus 2002; Christensen et al. 2016; Fridland 2017; Wu 2020; Shepherd 2021; Pavese 2021a).

Besides control, the other characteristic feature of skills such as archery, playing the piano, gymnastics, carpentry is that they are *learnable* (Ryle 1949; Singleton, 1978; Pavese forthcoming). Their learnability puts them in opposition to *instincts*. Instincts are

inherited and unlearned abilities for even fairly complex and stereotyped activities which are common to members of a species (Blumberg 2017). Instincts are not learned; they are *developed*. Philosophers of biology distinguish an instinct's development from a skill's learnability in terms of whether it develops *across a variety of different environments*. A skill exhibits a modal dependence on environments: *had the environments been different, the skill would have been different too* (cf. Stich 1975). The general idea is that, while an instinct tends to develop in pretty much every healthy member of the species that grows in a normal environment, skills are only acquired by individuals who find themselves in appropriate environments. The acquisition of skills might depend on the resources made available by the environment—e.g., sailing skills develop in populations in proximity to water, whereas mountaineering skills develop in populations living in mountainous regions. The differing availability of materials makes for widely different tool-use skills—e.g., Incas' tool skills were shaped by the availability of stone, copper, and bronze, but not of iron. Relatedly, the acquisition of skills is at least partly dependent on the social environment. Skills are faster to acquire by imitation and verbal feedback—by social learning. Thus social learning considerably facilitates the acquisition and transmissibility of skills (Hosfield 2009). By contrast, instincts are generally thought to be capable of developing in socially deprived environments—i.e., even in socially isolated animals (Lorenz 1981). Finally, skills are learnable to the expert level only through active, or deliberate practice (Ericsson 2008)—one can become an expert only by practicing with the intention of improving and fixing mistakes; rote repetition will not do. These processes of learning are rational—involves evidence gathering of the mistakes

one tends to make, hypothesis testing of better ways of performing the task, and subsequent rational revisions. By contrast, mere habits can be learned without such rational processes—rote repetition and reinforcement learning suffice for acquiring a habit.

Paradigmatic examples of skills such as archery, chess, basketball, playing a musical instrument, carpentry or tool use also differ in some crucial epistemological respects also from *general faculties*, such as memory, perception, attention, reasoning, or imagination. Although the general faculties can sensibly improve by practice and training, they differ from skills in that general faculties are mental abilities that we cannot but develop—as humans we cannot fail to eventually develop, to at least some degree, memory, imagination, reasoning, attention, perception.

This ‘narrow’ notion of skills that I have isolated comes with a long pedigree. Indeed, it arguably traces back to Aristotle’s conception of *technē*. Aristotle did not count general faculties such as perception, memory, reasoning or imagination among the *technai*. Paradigmatic examples of Aristotelian *technai* include, crafts, such as the art of building (cf. *Metaphysics*. E.2 1026b6-10), carpentry, knitting, or tool use, but also medicine and the art of war (*Posterior Analytics*. II.11 94a36-94b8), intellectual skills such as rhetoric, as well as sport skills, such as gymnastics and sailing. One of key texts for Aristotle’s discussion of *technē* is *Metaphysics* Θ.2. The passage distinguishes between powers in accordance with *logos* (*meta logou dunameis*) and powers without *logos* (*alogoi dunameis*). *Technai* are powers in accordance with *logos* (cf Johansen 2017). He also tells us that *technai* are productive forms of knowledge (*poiētikai*

epistēmai), the manifestation of which are *products*—the effects of particular kinds of actions, or productions. Aristotle makes it clear that an artisan is able to teach their *technē* (*Eudemian Ethics* I.8 1218b16-22; *Metaphysics* E.2 1027a20-23, *Topics* VI.4 141a29-30), since *technai* involve causal explanatory accounts (*logoi*) that can be taught and so can be transmitted to others. Aristotle also tells us that *technai* are two-way powers, or *powers for contraries*.

In this way, Aristotle distinguishes *technē* from other rational powers—such as virtue, which is a one-way power. The contrast between *technē* and *virtue* is at the core of Aristotle's discussion in the *Ethics*, where he tells us that "... in a craft, someone who makes errors voluntarily is more choiceworthy; but with prudence, as with the virtues, the reverse is true." (*Nicomachean Ethics* VI. 5 1140b20-25). The crucial distinction seems to be one in *voluntary control*. One way of making an error voluntarily is *to refrain from exercising the power while still having it*. For instance, the housebuilder can be skilled at their job while refraining in some cases to exercise their art: their refusing a commission does not undermine their status as skilled housebuilder. Moreover, there is no difficulty, for a skilled agent, in knowingly disregarding what it is right to do in light of the norms of a certain skill-domain. For instance, while playing chess, one might recognize that castling is the best way to avoid being checkmated but might choose to make some other move instead if they desire to lose. This choice does not automatically call into question one's skill as a chess player. Virtues, by contrast, are one-way powers, since a person cannot refrain from acting generously when required while still counting as generous. When it comes to virtue, *ceteris paribus*, someone's failure to act kindly in relevant

situations immediately calls into question whether she is rightly credited with possessing the virtue of kindness. This idea that skills differ from virtues in that skilled agents retain a control over whether or not to exercise or refrain from exercising a given skill is rather standard among contemporary philosophical literature on the virtues.

Now consider the relation between skills and epistemic competences. Our discussion has highlighted one crucial dimension of difference: like general faculties, epistemic competences are not learnable in the same robust sense in which skills are learnable. It is a vexed question whether the characteristic manifestations of epistemic competences (beliefs, or knowledge) are action-like. Some sharply distinguish between beliefs and knowledge (states) on the one hand, and actions on the other (Engel 2013), though this distinction is challenged by examples such as marriage and friendship—i.e., states which require activity for maintenance. So perhaps beliefs too can be conceived of on the model of activities (Rohrbaugh 2015). Indeed, perhaps the sort of modal robustness that many epistemologies require of beliefs to count as knowledge is just a special instance of the sort of modal robustness that is required of skilled action (Beddor and Pavese 2020). Finally, a question concerning epistemic competences is whether they are more akin to either skills or virtues when it comes to their being *voluntarily controlled*. Exercising one's epistemic competence is typically a matter of forming beliefs in response to considerations that bear on what it is correct to believe in light of the relevant epistemic norm. In this respect, epistemic competences appear more similar to virtues than to skills (Horst 2021).

Finally, consider the relation between skills and know-how. Artemisia Gentileschi could not be as skilled at painting without knowing how to paint. On the other hand, know-how might not suffice for skill possession. For example, it is undeniable that many mediocre painters knew how to paint but not sufficiently well to count as skilled. On the other hand, while knowing how to paint does not entail skill, knowing how to paint sufficiently well might. Finally, one might think some know-how is innate, so it might not involve robust learning in the same way as the notion of skill that we have discussed. It is certainly true that, while early philosophical literature at the boundaries between action and epistemology has focused on know-how at the expense of skills, more recently the literature has evolved to draw a clearer distinction between the topics of know-how and skills and to explore skills in their own right in further detail.

2. Puzzles about learning

How do people learn the skill to, say, play the piano? Say that at t , a subject S does not have the skill and at a later time, S acquires the skill, at the highest level of expertise.

What must have happened between time t and time $t+n$ in order for S to have acquired the skill to play the piano? Simple as this question apparently is, many tricky issues arise when one tries to answer it.

Current psychological theories tell us that learning a skill at a proficient level takes on average around ten thousand hours of practice (Ericsson 2008). Thus, presumably, between t and $t+n$, S must practice to play the piano. What does practicing a skill involve? Plausibly, for a skill to ϕ , practice must require ϕ -ing a bunch of times. Now,

suppose one does need to ϕ intentionally in order to practice ϕ -ing. For example, one must play the piano intentionally a bunch of times in order to practice playing the piano. This supposition, though *prima facie* plausible, raises an issue: intentionally ϕ itself requires some degree of skill. If a novice at darts, with no skill whatsoever, hits the bullseye on their first try, their success does not seem intentional (Malle and Knobe 1997; Pavese et al 2023). Some degree of skills is usually thought to be required for the sort of control in action that intentionally ϕ necessitates (Knobe 2003; Pavese and Beddor 2022). Moreover, only action that is skilled to some degree is creditable: action that is not due to skills at all is due to luck, and so cannot be credited to its agent. If so, with no degree of skill whatsoever, one cannot intentionally ϕ at all—one can only ϕ non-intentionally.

These weak assumptions are sufficient to raise a puzzle from learning. By assumption, at t , when S starts practicing for the very first time, S does not possess any degree of skill at all. If some degree of skills is needed for intentionally ϕ , S cannot intentionally ϕ at t . If practice to ϕ requires intentionally ϕ -ing, then it seems S cannot even start practicing before acquiring some degree of skill. Yet, plausibly, in order to acquire some degree of skill, one does need to practice to ϕ . Thus, we got a puzzle. We can put the quandary a bit more schematically as follows:

Learning a skill: Acquiring the skill to ϕ requires practicing to ϕ .

Practice: Practicing to ϕ requires ϕ intentionally.

Intentionally: Intentionally ϕ requires some degrees of skill to ϕ .¹

¹ Pineros Glasscock (2022) also raises a puzzle from practice but there are doubts as to whether his puzzle is genuine. See Pavese forthcoming for discussion.

From **Learning a skill**, **Practice**, and **Intentionally**, it seems to follow that one cannot acquire the skill to ϕ by practicing to ϕ , since by **Practice**, practicing to ϕ requires intentionally ϕ . But at the very beginning, one has no degree of skills at all. So, by **Intentionally**, one cannot intentionally ϕ , and hence cannot even start practicing to ϕ . Thus, it seems that there is no way for S to learn the skill, at t , if one does not already have at t the skill to ϕ . Hence the paradox.

Abandoning **Intentionally** would leave us with no plausible control constraint on intentional action (Pavese and Henne 2023). On the other hand, **Learning a skill** seems tautological. So if we are to get out of the puzzle of learning by practice, we are left rejecting **Practice**. Rejecting **Practice** amounts to saying that practicing to ϕ does not necessarily require intentionally ϕ . How could that be so? The thought is that, especially at the very beginning of the learning process, when one does not have the skill to ϕ to any degree, practicing to ϕ might not require to intentionally ϕ *already*. Consider replacing **Practice** with the weakened:

Practice*: Practicing to ϕ requires only attempting to ϕ .

The problem though is that **Practice*** seems to underwrite an implausible view of practice: if practicing to ϕ only required attempting to ϕ and not actually to ϕ , then it seems that by practice one could get better only at attempting to ϕ , without actually getting better at ϕ -ing. A better option might be to weaken **Practice** into **Practice***:

Practice**: Practicing to ϕ requires ϕ -ing.

On this view of practice, practicing to ϕ does not necessarily require intentionally ϕ -ing, especially at the start of the learning process, where one is not in position to intentionally

ϕ , precisely due to lack of skills. By replacing **Practice** with **Practice****, we get out of the puzzle, since ϕ -ing, when ϕ is not intentional, does not require any degree of skill.

Interestingly, the puzzle from learning resurfaces for essentially intentional actions. According to an influential doctrine, according to which certain actions are special, in that they can only be performed intentionally (Anscombe 1957, §47). Some putative examples are actions such as *greeting* and *marrying*. Let ψ be an essentially intentional action. Then by **Acquiring a Skill**, we have that acquiring the skill to ψ requires practicing to ψ . By **Practice***, practicing ψ requires ψ ; but if ψ is essentially intentional, then practicing ψ requires ψ intentionally. By **Intentionally**, ψ intentionally requires some degree of skills. But at t , when starting to learn to ψ , S has no degree of skills, by assumption. Hence, we got a puzzle back again. Some have independently questioned the doctrine of essentially intentional actions (Beddor and Pavese 2022). This puzzle from learning seems to offer one more reason to abandon it (Pavese forthcoming).

2.1 Regresses

Ryle (1946, 1949) oscillates between a broader conception of skills that includes general faculties and competences and a narrow conception of skills conceived of as *technai*. For example, Ryle (1946: 7) talks of reasoning—a general faculty—as a skill. According to Kremer (2020), Ryle would also include perceptual abilities and epistemic competences among skills. Indeed, insofar as knowledge is formed through perception and other epistemic competences, Ryle thought that propositional knowledge was an activity that too required skill. Yet, Ryle also embraces the claim that a skill's characteristic

manifestations are necessarily controlled, as when he (1949: 5) compares the skilled clown—who tumbles on purpose—to the klutz, who only does so unintentionally. Ryle talks of skills as a distinctive form of knowledge, inculcated through a distinctive form of teaching or ‘training’ (Ryle 1967). And he does seem to think of training in terms of what we have characterized as robust learning. Indeed, Ryle (1949: 42-3) distinguishes habits from skills in that, although both are *acquired dispositions*—the former are acquired through ‘drill,’ ‘conditioning,’ or ‘mindless repetitio’, whereas the latter are acquired through *training*, which involves ‘the stipulation by criticism and by example of a pupil’s own judgment,’ in which the pupil ‘learns how to do things thinking what he is doing’. Finally, for Ryle, skills are know-how, which he thought of as a particular kind of intelligent disposition—a ‘multi-track’ disposition. One distinctive feature of knowing how to *F* is that this knowledge does not need to be exercised in acts of *F*-ing. For example, Ryle thought that a chess player might manifest their know-how and skill simply by recognizing what the right move is, even if they refrain from making it.

Ryle’s chief argument for thinking that skills can neither be nor involve propositional knowledge is the well-known *regress argument*. In the original formulation, (Ryle 1949, 19-20), the argument assumes that skillful behavior involves contemplating knowledge—as ‘intellectualist legend’ has it. Contemplating knowledge is itself behavior that can be skillful or unskillful. If this particular skillful behavior also requires some further act of contemplation, a regress emerges. While this regress assumes the objectionable assumption that using knowledge requires contemplating or consciously

entertaining some proposition, the same regress can be reformulated on the weaker assumption that skillful behavior requires *using* knowledge (Pavese 2021c).

The regress seems to require, for anybody to act intelligently, that one ought to perform an infinite number of acts *of using knowledge*. And one might worry that this is temporally and psychologically impossible—because any such act will require some time to perform, thereby making it impossible to perform any intelligent action. And yet, this cannot be what makes the regress vicious. After all, a temporal construal of the regress makes it analogous to Zeno’s paradox, of which there are well-known solutions. Nevertheless, the regress argument commits us to something false, regardless of whether or not the regress commits us to something paradoxical. As stated, the regress commits us to the claim that the intelligence of an act depends not only on whether that act uses knowledge, but also on whether an infinite number of acts use knowledge, thereby making the question of whether any act is intelligent entirely extrinsic to that action. One might be forgiven for wanting a theory of skillfulness that does not require this unexpected commitment.

What view did Ryle embrace as an alternative to the intellectualist legend?

Though this is sometimes debated, Ryle must have thought that skills and know-how are dispositions that are not grounded, not even in part, in propositional knowledge at all. Skills cannot require any contentful state such as knowledge for Ryle, since this claim, together with another that he seems to have endorsed, would quickly lead his view to another kind of regress. As we have seen, Ryle thought of knowing as involving the activity of grasping a proposition, which requires the exercise of epistemic skills if it has

to yield knowledge. However, if skills required knowledge, since knowledge does require skills, Ryle's dispositionalism would also quickly lead to a vicious regress—according to which, in order to know how to perform *any task at all*, one would need to perform an infinite number of graspings. Indeed, any view that takes skills to constitutively involve knowledge, while falling short of identifying skill with them, *already* counts as intellectualist by Ryle's own lights, since any such a view is already exposed to a version of Ryle's original regress argument.

2.2 Old Intellectualisms

In their defense of the Intellectualist legend, Stanley and Williamson (2001) focused on know-how, rather than on skills. According to them, knowing how to *F* is a matter of knowing a proposition, but this knowledge does not need to endow one with an ability—not even with a general ability, or skill. So, for them, knowing how is just a kind of knowledge-wh. Their primary argument for this view of knowing how relies on the semantics of the ascriptions of knowledge-how and the analogy with other sorts of ascriptions of propositional knowledge. This argument faces the objection that it narrowly focuses on English ascriptions, since in other languages the analogy between ascriptions of knowledge-how and propositional knowledge is much looser (e.g., Rumfitt 2003; Brown 2013).

Stanley and Williamson (2001) developed a response to Ryle's regress argument that Ginet (1975) had originally sketched. Ginet (1975) questioned the assumption that at every point in the chain of the regress, using knowledge can be assessed for skillfulness,

since he objected that only intentional actions are the sort of things that can be skillful or unskillful. Non-intentional behavior, such as digesting one's food, is not the sort of thing that can be skillful or unskillful. The idea behind this response is that in order to stop the regress chain, there will be operations that amount to using knowledge which cannot themselves be assessed for skillfulness because they are not intentional (cf. Fodor 1968 for a similar response).

Though this response has been the main defensive strategy adopted by intellectualists so far, it is not entirely unproblematic. Consider remembering to check the rearview mirror when driving. This behavior need not be intentional. And yet, if one's aim is to reach one's destination safely, remembering to check the rearview mirror can be intelligent, albeit not intentional (Weatherson 2017).

One might insist on denying that behavior can be skillful, if not intentional. This move might seem appealing at first. For example, consider unintentionally offending one's boss. Though the offense might not be intentional, it plausibly consists in one intentionally performing some actions—such as saying or insinuating certain things in front of one's boss. Or consider remembering to check one's rearview mirror or remembering to turn the gas off. One might contend that it is not so much that we assess the act of remembering for intelligence *per se*; rather, we assess the intelligence of the actions that this remembering gives rise to—such as the checking of the rearview mirror, or the turning off of the gas.

The problem with this line of response is that it does not clearly cover all possible counterexamples. Consider *forgetting* to check the rearview mirror. This is not skillful.

And yet there is no set of intentional actions to which forgetting to check the rearview mirror corresponds. Another way of blocking the regress argument is to offer reason for thinking that the culprit of the regress ought to be relocated (Pavese forthcoming). The reason is that the regress arises even *without* assuming intellectualism. One gets the regress up and running already on the assumption that skillful behavior uses one's skill or knowledge-how. For using knowledge-how is behavior that can itself be assessed for skillfulness (i.e., that can be skillful or unskillful). Just as before, this behavior will be skillful only if it itself uses knowledge-how; this use of knowledge-how in turn, if skillful, will use knowledge-how, and so on. If the regress arises without assuming that there is a constitutive connection between skillfulness and knowledge, then the culprit of the regress cannot be this assumption. So, which is the culprit?

All extant versions of the regress assume that intelligent behavior is skillful behavior. Since intellectualism claims that skillful behavior uses knowledge, then it follows that intelligent behavior must use knowledge. But there are independent reasons to challenge the claim that intelligent behavior is skillful. Indeed, this assumption gives rise to a different kind of regress—the *learning regress*: assuming every intelligent behavior is learned, how can we account for the intelligence of the very first acts of learning? Some suggest that the intelligence of the learning inherits its skillfulness from the teacher of the skill (Small 2014). The problem with this reply is that skills do not *always* need to be learned socially—individual learning will often do. If individual learning episodes can be intelligent, then there must be intelligent behavior that is not skillful, since it is not itself learned. On the other hand, if some intelligent behavior is not

skillful, one can be an intellectualist in maintaining that skillful behavior uses knowledge, and blocking the regress by allowing that in the chain of acts, some are intelligent without being skillful, and so without using knowledge.

Whether knowledge-how could be knowledge-that has been attacked from the epistemological point of view. Cath (2011) argues that knowledge-how resists Gettierization, and so do Carter and Pritchard (2015). Their arguments rely on intuitions concerning knowledge-how ascriptions in Gettier cases. Recent evidence from experimental philosophy shows that these conclusions might be too quick. Carter, Shepherd, and Pritchard (2019) find that in environmental luck cases, people are as inclined to ascribe knowledge-that as they are to ascribe knowledge-how. Pavese, Henne, and Beddor (2023) have gathered evidence that in classic Gettier cases, people actually refrain from ascribing knowledge-how as much as they refrain from ascribing knowledge-that.

2.3 Epistemic Dispositionalism

According to epistemic dispositionalism, skills are dispositions to know. In particular, one is skilled at φ -ing if and only if S is disposed to have knowledge that is appropriate for guiding tokens of φ -ing (Löwenstein 2017; Habgood-Coote 2019; Stanley & Williamson 2017).

One first issue for epistemic dispositionalism is that it obliterates the epistemological differences between various epistemic competences such as perception and skills, such as archery, chess, and playing the piano. Perception is arguably a

competence to know (Millar 2010). So the view cannot account for the robust learnability that we have seen is distinctive of skills as opposed to epistemic competences. By embracing epistemic dispositionalism for skills but intellectualism for know-how, Stanley and Williamson (2017) divorce skills from know-how. And yet one would have thought that, if a subject is skilled at a task, then they know how to perform it.

Epistemic Dispositionalism is in trouble when it comes to underwriting the connection between skillful action and controlled action (Pavese and Beddor 2022). This is the case because dispositions to know characteristically manifest in knowledge states, not actions, so not the sort of things that can be controlled. Perhaps, then, a better proposal on behalf of epistemic dispositionalism is that the only *secondary* manifestations of skills qualify as skillful. Now, Stanley and Williamson do argue that skills secondarily manifest in actions guided by knowledge states (2017: 717). While this avoids obvious counterexamples, there is a worry that the definition of secondary manifestation here is arbitrary. After all, we would like our definition of secondary manifestation to follow from a *more general* theory of disposition manifestation—one that is not itself tailored to skillful action. A natural generalization would be that any disposition *D* secondarily manifests in whatever states or actions are explained by the primary manifestation of *D*. Note, however, that this generalized definition of secondary manifestation does not predict that only skillful actions are intentional. For example, if an athlete's prowess induces envy in an onlooker, the envy is explained by the primary manifestation of their athletic dispositions. But the envy is neither intentional nor skillful. So, regardless of whether or not we focus on primary or secondary manifestations,

epistemic dispositionalists fail to predict that only controlled actions are skillful. Thus, epistemic dispositionalism fails to capture what is distinctive of skills: robust learnability and control. As such, epistemic dispositionalism might turn out to be a better view of epistemic competences than of skills proper (Kelp 2018).

2.4. Cognitivism about skills

One argument in favor of cognitivism about skills—according to which knowledge enters constitutively in skills—can actually be traced back to Aristotle. In general terms, Aristotle’s argument for thinking that *technai* involves knowledge is premised on the assumption—which we have already discussed—that *technai* are two-way powers, or powers for contraries. Aristotle’s main argument for thinking that *technai* are powers for contraries is that productive knowledge is of contraries (10464ff). This conclusion follows only if *technai* are productive forms of knowledge. Later in *Metaphysics* Θ.2., Aristotle says that productive knowledge is an account (*logos*). And it is clear from his examples and other passages that the account in question is an account of the product of the relevant *technē*. In *Metaphysics* Z.7, Aristotle further confirms that a technical *logos* is an account of what the product is, namely its essence (*to ti ēn einai*, 1032b1-2; cf. *Physics* II.1 193a30-36). To have a *technē* is to have a form of the product in the soul, where the form is what something is. That is, artisans in some sense grasp the essence of the products they make. Thus, doctors know what health is, whereas builders know what a house is. This raises the question of what exactly the account of the product looks like. We know from other passages that for Aristotle such accounts concern *universals* and are

explanatory. *Metaphysics* A.1 contrasts people with *technē* with people who merely have experience (*empeiria*, cf. *Metaphysics* A.1 981a10-25). The merely experienced person can tell that a particular person will benefit from a given cure since *empeiria* is concerned with particulars (981a5-24, esp. 15-16). But the doctor can also tell, for instance, what kind of cure benefits a given class of people—say, that such and such cure benefits phlegmatic people with fevers (981a7-12). So, *technē* is associated with knowledge which consists of generalizations that exclusively belong to the level of universals. People with mere experience only know that something is the case—say, that giving a certain medicine cures fever. But people with *technē* can explain *why* a given medicine is the right cure for fever (*Metaphysics* A.1 981a25) because they “know the cause” (981a30). Here to know the answer to ‘Why X?’ is tantamount to knowing the cause of X, where ‘the cause’ picks one of the four Aristotelian causes (cf. *Physics* II.3 194b18-20, *inter alia*). So, one can equally say that technical accounts are *causal*, in that they individuate the cause(s) of their corresponding products.

According to Aristotle, *technai* are (a) for contraries since productive knowledge is for contraries and (b) possessing productive knowledge of a product X entails possessing productive knowledge of the contrary of X. More specifically, given that technical knowledge is causal-explanatory, one is in a position to grasp an explanation of the contrary X in virtue of grasping explanations of X. For if one grasps explanations of X, and knows that Y is the absence of X in a given type of entity, one is in a position to explain the absence of X as well. Thus, Aristotle’s argument for intellectualism (for thinking that *technai* must be explanatory knowledge) is premised on the claim that

technai are two-ways powers and hinges on the kind of productive knowledge that the possessor of *technê* must have in order to possess a two ways power.

Cognitivism has similar resources to develop an argument for thinking of skills as knowledge states that starts both from control and from robust active learnability. According to a prominent tradition that traces back to Anscombe (1957), skilled action requires knowing what one is doing when doing it—knowing what one is doing captures the sort of control that is distinctive of skilled action. Many agree that knowing what one is doing is a matter of knowing which means to take to which ends when acting skillfully (Setiya 2012; Pavese 2022).

Suppose control requires practical knowledge. This by itself does not mean that it requires a standing propositional knowledge state, since practical knowledge is just an occurring knowledge state. In order to make a second step towards neo-intellectualism, consider Gentileschi skillfully painting her *Self-portrait as the Allegory of Painting* (1638-1639). Her control requires her to know, when painting, which means to take to which ends. What guarantees that her practical knowledge guides her execution? It would be hard for anybody to know what one is doing when executing an action that one does not know how to execute. This observation suggests that practical knowledge requires knowledge-how. While practical knowledge is an occurring state of knowledge that explains control in action, knowledge-how is a standing knowledge state. So, control of the sort that is exhibited in skilled action demands that one knows how to perform it ahead of performing it—it requires a standing knowledge-how state.

How are we to think of such a knowledge-how state? Here comes the third crucial step on behalf of cognitivism. Recall that skills are learnable in the robust sense—they *can* be acquired through social learning, through imitation, practical, and verbal feedback. Indeed, the importance of teachability for *technē* was emphasized by Aristotle. According to Aristotle, not only is *technē* teachable. Aristotle famously endorsed a stronger claim—i.e., that the possessor of *technē* must be able to teach. For example, in *Metaphysics* A.1, just after clarifying that *technē* is more honorable than experience, Aristotle adds: “On general, being able to teach is a sign of whether a man knows or not, and because of this we consider *technē* to be more truly knowledge than experience is; for artisans can teach, but men of mere experience cannot” (981b7-9).

The claim, of course, is not that skills ought to be socially learned (Shepherd 2021)—in some cases, it might be possible to learn a skill, at least to a significant degree, by individual means alone (Pavese 2021a). The point, rather, is that skills are *susceptible* to being socially learned—through imitation and verbal feedback. The possibility of social learning is central to the acquisition and transmissibility of pretty much any skill that deserves the name. Morgan *et al.* (2015) argue that the teaching of general concepts such as that of a platform edge contributes to the development and transmissibility of Oldowan stone knapping techniques. The impact of verbal feedback has been shown to significantly affect the acquisition of wide ranging kinds of skills, from basic motor skills (e.g., Sullivan *et al.* 2008), to more complex sport skills such as swimming and tennis (e.g., Zatoń & Szczepan 2014), yoga skills, manual therapy skills, as well as surgical skills (e.g., Flinn *et al.* 2016) and musical skills (Duke & Henninger 1998), etc.

In order to model the impact of verbal feedback on the standing state of know-how, at least some of this learning ought to take the form of rational revision, evidence gathering, and hypothesis testing. Moreover, as we have seen, skills can develop into forms of expertise only through deliberate practice. But if skills are susceptible to be acquired in this fashion, with the intention of improving and fixing mistakes, then it seems that they ought to encompass a cognitive state that is updated and revised through imitation, practice, and verbal feedback (Pavese 2021d, forthcoming). If so, the distinctive learnability of skills—their active and social learnability—provides the final ingredient in an argument for the view that skills—as understood apart from habits, instincts, general faculties, and as corresponding to areas of expertise—constitutively involve a standing state of propositional knowledge.

References

- Anscombe, G. E. M. (1957). *Intention*, Oxford: Basil Blackwell.
- Beddor, B. and Pavese, C. (2020). “Modal Virtue Epistemology”, in *Philosophy and Phenomenological Research*, 101(1): 61–79.
- Beddor, B., & Pavese, C. (2022). “Practical knowledge without luminosity.” *Mind*, 131(523), 919-936.
- Brown, J. A. (2013). “Knowing-How: Linguistics and Cognitive Science”, *Analysis*, 73(2): 220–227.

- Blumberg, M. S. (2017). “Development evolving: the origins and meanings of instinct.” *Wiley Interdisciplinary Reviews: Cognitive Science*, 8(1-2), e1371.
- Carter, J. A. and D. Pritchard (2015). “Knowledge-How and Epistemic Luck”, *Noûs*, 49(3): 440–453.
- Carter, J. Adam, D. Pritchard, and J. Shepherd (2019). “Knowledge-How, Understanding-Why and Epistemic Luck: An Experimental Study”, *Review of Philosophy and Psychology*, 10(4): 701–734.
- Cath, Y. (2011). “Knowing How Without Knowing That”, in Bengson and Moffett 2011c: 113–135.
- Christensen, W., J. Sutton, and D. McIlwain (2016). “Cognition in skilled action: Meshed control and the varieties of skill experience”, *Mind & Language*, 31(1): 37–66.
- Davidson, D. (1971). “Agency”, in Robert Binkley, Richard Bronaugh, and Ausonia Marras (eds.), *Agent, Action, and Reason*, Toronto: University of Toronto Press, reprinted in Davidson 2001a.
- Dreyfus, H. L. (2002). “Intelligence without Representation—Merleau-Ponty’s Critique of Mental Representation: The Relevance of Phenomenology to Scientific Explanation”, *Phenomenology and the Cognitive Sciences*, 1(4): 367–383.
- Dreyfus, S. E. (2004). “The five-stage model of adult skill acquisition.” *Bulletin of science, technology & society*, 24(3), 177–181.

Duke, R. and Henninger, J. (1998), “Effects of verbal corrections on student attitude and performance” in *Journal of Research in Music Education* 46(4).

Engel, P. (2013). “Sosa on the normativity of belief.” *Philosophical Studies*, 166(3), 617–624.

Ericsson, A. K. (2008). “Deliberate practice and acquisition of expert performance: a general overview.” *Academic emergency medicine*, 15(11), 988-994.

Fantl, J. (2011). “Ryle’s regress defended.” *Philosophical Studies*, 156(1), 121-130.

Fitts, P. M. and M. I., Posner, 1967, *Human Performance*, Belmont, CA: Brooks/Cole.

Fodor, J. A. (1968). “The Appeal to Tacit Knowledge in Psychological Explanation”, *The Journal of Philosophy*, 65(20): 627.

Fridland, E. (2013). “Problems with Intellectualism”, *Philosophical Studies*, 165(3): 879–891.

Fridland, E. (2017). “Skill and Motor Control: Intelligence All the Way Down”, *Philosophical Studies*, 174(6): 1539–1560.

Fridland, E. and C. Pavese (eds.) (2020). *The Routledge Handbook of Philosophy of Skill And Expertise*, New York: Routledge.

Gettier, E. L. (1963). “Is Justified True Belief Knowledge?”, *Analysis*, 23(6): 121–123.

Ginet, C. (1975). *Knowledge, Perception and Memory*, Dordrecht: Springer Netherlands.

Goldwasser, S. (2023). "Memory as skill." *Review of Philosophy and Psychology*, 14(3), 833-856.

Habgood-Coote, J. (2019). "Knowledge-How, Abilities, and Questions", *Australasian Journal of Philosophy*, 97(1): 86–104.

Heatherington, S. (2020). "Knowledge as Skill", in Fridland and Pavese 2020: 168–178.

Horst, D. (2021). "Is Epistemic Competence a Skill?" *Australasian Journal of Philosophy*, 1–15.

Hosfield, R. T. (2009). *Modes of transmission and material culture patterns in craft skills* (Vol. 2, pp. 45-60). University of California Press.

Johansen, T. K. (2017). "Aristotle on the Logos of the Craftsman", *Phronesis*, 62(2): 97–135.

Kelp, C. (2018). *Good thinking: A knowledge first virtue epistemology*. Routledge.

Kind, A. (2020). "The Skill of Imagination." In *The Routledge Handbook of Philosophy of Skill and Expertise* (pp. 335-346). Routledge.

Knobe, J. (2003). "Intentional Action and Side Effects in Ordinary Language." *Analysis* 63, 190–193.

Kremer, M. (2020). "Gilbert Ryle on Skill as Knowledge-How", in Fridland and Pavese 2020: 100–112.

Lorenz, K. (1981). *The Foundations of Ethology*. Springer Science & Business Media.

Löwenstein, D. (2016). *Know-How as Competence: A Rylean Responsibilist Account*, Verlag Vittorio Klostermann.

Malle, B. F., & Knobe, J. (1997). “The folk concept of intentionality.” *Journal of experimental social psychology*, 33(2), 101-121.

Millar, A. (2009). “What Is It That Cognitive Abilities Are Abilities to Do?”, *Acta Analytica*, 24(4): 223–236.

Miller, E. K. (2000). “The prefrontal cortex and cognitive control.” *Nature reviews neuroscience*, 1(1), 59-65.

Morgan, T. et al. (2015). “Experimental evidence for the co-evolution of hominin tool-making teaching and language” in *Nature communications* 6(1).

Pavese, C. (2013). *The Unity and Scope of Knowledge*, Ph.D. Thesis, New Brunswick, NJ: Rutgers University.

Pavese, C. (2015a). “Practical senses” in *Philosophers’ Imprint*.

Pavese, C. (2015b). “Knowing a rule” in *Philosophical Issues*.

Pavese, C. (2016). “Skills in Epistemology I and II” *Philosophy Compass*: 11: 642–660.

Pavese, C. (2017). “Know-How and Gradability”, *Philosophical Review*, 126(3): 345–383.

- Pavese, C. (2019). “The Psychological Reality of Practical Representation”, *Philosophical Psychology*, 32(5): 784–821.
- Pavese, C. (2021a). “Knowledge and Mentality”, *Philosophical Perspectives* 35(1): 359–383.
- Pavese, C. (2021b). “Know-How, Action, and Luck”, in *Knowledge and Justification, New Perspectives*, special issue of *Synthese*, 198(7), pp. 1595–1617 2021.
- Pavese, C. (2021c). “Knowledge-How”, *Stanford Encyclopedia of Philosophy*.
- Pavese, C. (2021d). “Review of Shepherd’s The Shape of Agency”, in *Mind*, pp. 1–9.
- Pavese, C. (2022). “Practical Knowledge First”, *Synthese*.
- Pavese, C. (2024). “Intelligence Socialism ”, forthcoming in *Oxford Studies in the Philosophy of Mind*.
- Pavese, C. forthcoming. *The Practical Mind: Skills, Knowledge, and Intelligence*, Cambridge University Press.
- Pavese, C. & B. Beddor (2022) “Skills as Knowledge”, *Australasian Journal of Philosophy*.
- Pavese, C., P. Henne, and B. Beddor (2023) “Epistemic Luck, Knowledge-How, and Intentional Action” *Ergo*, 10: 36.

- Pavese, C., & Henne, P. (2023). “The know-how solution to Kraemer's puzzle.” *Cognition*, 238, 105490.
- Piñeros Glasscock, J. S. (2022). “The puzzle of learning by doing and the gradability of knowledge-how.” *Philosophy and Phenomenological Research*, 105(3), 619-637.
- Piñeros Glasscock, J. S. (2020). “Practical knowledge and luminosity.” *Mind*, 129(516), 1237-1267.
- Rosenbaum, D. A., Carlson, R. A., & Gilmore, R. O. (2001). “Acquisition of intellectual and perceptual-motor skills,” *Annual Review of Psychology* 52(1), 453–470.
- Rohrbaugh, G. (2015). “Inner achievement,” *Erkenntnis*, 80(6), 1191–1204.
- Ryle, G. (1946). “Knowing How and Knowing That: The Presidential Address,” *Proceedings of the Aristotelian Society*, 46(1): 1–16.
- Ryle, G. (1949). *The Concept of Mind*, University of Chicago Press.
- Ryle, G. (1967) [2009]. “Teaching and Training”, in *The Concept of Education*, R. S. Peters (ed.), London: Routledge, ch. 7.
- Rumfitt, I. (2003) “Savoir faire.” *The Journal of Philosophy*, 100(3), 158-166.
- Schmidt, R. A., & Wrisberg, C. A. (2008). “Motor learning and performance: A situation-based learning approach.” *Human Kinetics*.
- Setiya, K. (2012). “Knowing How”, *Proceedings of the Aristotelian Society*, 112(3): 285–307.
- Shepherd, J. (2021). *The shape of agency: Control, action, skill, knowledge*, Oxford University Press.

- Siegel, S. (2020). “Skill and expertise in perception”, in Fridland and Pavese (2020): 303–16.
- Singleton, W. T. (1978). “The analysis of practical skills.” Springer Science & Business Media.
- Small, W. (2014). “The transmission of skill.” *Philosophical Topics*, 85-111.
- Sosa, E. and L. F. Callahan (2020). “Skill and Knowledge”, in Fridland and Pavese (2020): 146–156.
- Stanley, J. and T. Williamson (2001). “Knowing How”, *Journal of Philosophy*, 98(8): 411–444.
- Stanley, J. and T. Williamson (2017). “Skill”, *Noûs*, 51(4): 713–726.
- Stich, S. P. (Ed.). (1975). *Innate ideas* (Vol. 10). Univ of California Press.
- Stichter, M. (2018). *The Skillfulness of Virtue: Improving our Moral and Epistemic Lives*, Cambridge University Press.
- Sullivan, K. et al. (2008). “Motor learning in children: feedback effects on skill acquisition” *Physical therapy* 88(6).
- Weatherson, B. (2017). “Intellectual Skill and the Rylean Regress”, *The Philosophical Quarterly*, 67(267): 370–386.
- Williamson, T. (2000). *Knowledge and Its Limits*, Oxford University Press.
- Wu, W. (2020). “Automaticity, Control, and Attention in Skill”, in Fridland and Pavese (2020): 207–218.

Zatoń, K. and Szczepan, S. (2014). “The impact of immediate verbal feedback on the improvement of swimming technique” in *Journal of Human Kinetics* 41.