Philosophical Studies
https://doi.org/10.1007/s11098-024-02261-0

Knowledge, skills, and creditability

Carlotta Pavese[1]

Accepted: 23 November 2024

Abstract
The article discusses the relation between skills (or competences), creditability, and aptness. The positive suggestion is that we might make progress understanding the relation between creditability and aptness by inquiring more generally about how different kinds of competences and their exercise might underwrite allocation of credit. Whether or not a competence is acquired and whether or not a competence is actively exercised might matter for the credit that the agent deserves for the exercise of that competence. A fine-grained taxonomy of competences opens up the possibility of instinctual knowledge (knowledge by mere instincts) as well as the possibility of habitual knowledge (knowledge by mere habits), alongside knowledge by skills (or alongside knowledge by yet other sorts of competences). If instinctual knowledge were possible, it is suggested that it might not be of the sort that deserves credit at all. By piggybacking from the literature in evolutionary psychology, I suggest that, as inborn social learners, merely instinctual—and so not fully creditable—knowledge might be a reality for us.

Keywords Skills · Virtue epistemology · Competences · Aptness

1 Introduction

Sosa's (2007, 2011, 2015, 2021) virtue epistemology offers an elegant axiological framework that applies both to the epistemological domain and to the practical domain. In looking at the epistemic domain from the more general point of view of performances and activities, it promises to shed light on the normative parallels between the epistemic and the practical. Crucial to Sosa's unified axiological framework is understanding knowledge in terms of a property—that of *aptness*—that is supposed to apply to both practical and epistemic performances. Aptness is often elucidated by Sosa by invoking the concept of "creditability:" apt performance is fully creditable to its possessor (quotes to follow).

✉ Carlotta Pavese
carlotta.pavese@gmail.com; cp645@cornell.edu

[1] Cornell University, Ithaca, NY, USA

Published online: 10 December 2024

The topic of this article is precisely the normative status of aptness and in particular its relation to credit. Some critiques of Sosa's framework have targeted the very idea that apt belief, or knowledge, is creditable to their possessor (e.g., Lackey, 2009).[1] Though some of my discussion might be taken to add fuel to these critiques, my main goal is not to directly challenge the claim that knowledge is creditable. Rather, my main goal is explorative—i.e., I would like to explore the relationship between creditability and aptness, by locating it within wider explanatory contexts in action theory and in the philosophy of skills. My primary interest is in the questions: Why think that the creditability of a performance has to do with the exercise of skills? How to understand skills in such a way to underwrite that the claim that success in virtue of skills is fully creditable to its agent?

While my main goal is explorative, I will end up positively suggesting that a more fine-grained taxonomy of competences than the one proposed by Sosa—in particular a taxonomy that distinguishes between skills, instincts, faculties, and habits—might elicit different understandings of the relation between aptness and creditability. I will raise the possibility of merely instinctual knowledge and compare its credit status to that of apt (skilled) knowledge.

2 Aptness and creditability

Sosa defines "aptness" in terms of one's exercise of skills (or competences):

> **Aptness:** A performance is apt just in case it is successful in virtue of the exercise of a skill (or competence).

Aptness is Sosa's official view of aptness. Often, however, Sosa glosses "aptness" in terms of creditability. Here are six quotes—the last three from *Epistemic Explanations*, the most recent book by Sosa:

> Q1: A second proposal requires the belief to be **apt, correct in a way creditable to the believer,** as determined by how salient is the believer's competence in the explanation of his being right (Sosa, 2007: 80, my bold).

> Q2: A certain archer's shot hits the mark through a normal exercise of skill, let us suppose, in normal circumstances. [...] What if a gust of wind or stroke of lightning *might easily* have denied him his propitious situation for part at least of the relevant period, by affecting the arrow on its way to the target? Even so, the shot might have been **apt, surely, still accurate because adroit, and creditable to the agent,** so long as the competence remained in place, and the conditions appropriate, even if only by luck. What matters is that the conditions remain appropriately normal (or bet-

[1] Other virtue epistemologists underwrite the claim that knowledge is apt and creditable—e.g., Greco (2000, 2003), Zagzebski (1996), Riggs (2002), and Green (2016). Given the focus of this symposium, in this article, I will only discuss Sosa's views. For a version of virtue epistemology that does without aptness, see Beddor and Pavese (2020).

ter) along dimensions relevant to the agent's retained competence (Sosa, 2007:82, my bold).

Q3: **When a success, practical or intellectual, is creditable to an agent, it is due to an aptitude (to a competence or skill or virtue) seated in that agent**, whose exercise is rewarded with success in his act or attitude (Sosa, 2007: 87, my bold).

Q4: If you do not at all base your judgment on whether you have met them, then even if you opt successfully by affirming, **this is by luck in some measure, so that your success is not fully creditable to you** (Sosa, 2021: 64, my bold).

Q5: **A shot is apt and to the archer's credit** so long as the arrow is *in fact* unaffected by wind on the way to the target, *no matter how likely a spoiler gust may have been* (Sosa, 2021: 193, my bold).[2]

Q6: Consider our shot's pertinent spatiotemporal volume from the moment when the arrow is released to the moment when it strikes the target. Success in that volume would be quite unlikely despite our archer's excellent skill and shape. What makes success so unlikely is the very high risk (by hypothesis) of a spoiler gust. However, so long as no spoiler gust *in fact* comes along, our archer enjoys the complete competence required for **creditable, apt success** (Sosa, 2021: 193, my bold).

In Q2, Q5, Q6, Sosa treats "apt" as equivalent to "to an agent's credit" or to "creditable to the agent." Q1 concerns specifically doxastic performances, so "apt" is glossed as "creditable to the believer." In Q3 and Q4, Sosa commits to full creditability entailing aptness. These quotes together strongly suggest that Sosa endorses also:

Creditability: A performance is fully creditable to their agent just in case it is apt.

Now, of course, **Aptness** and **Creditability** are not equivalent. The former is a claim about the role of skills in apt performances. The latter is a claim about the relation between aptness of a performance and the credit that the performer deserves for it.

Indeed, **Creditability** is a substantive claim about aptness. Why should it be that only actions due to an agent's skills are fully creditable to that agent? What is it about skills and their exercise that make them well-suited to ground credit allocation? How to understand skills and competences in such a way to support the plausibility of **Creditability**?

[2] Other times, Sosa talks of credit not as a relation between an agent and a performance but as a relation between an agent's skill and the success of a performance (e.g., Sosa, 2021: 190). In my discussion, I will not focus on such uses of 'credit'.

Next section engages in a discussion of these issues. Preliminarily, let me dwell on the normative force of *being creditable* and on the importance of **Creditability** for Sosa's virtue epistemology.

I will take it that the creditability of a performance is a matter of whether the agent *deserves* credit for it. So understood, the force of creditability is that of a *ought*, not than of a *can*: a success is creditable to an agent if we ought to give the agent credit for it.[3]

So understood, **Creditability** is crucial to Sosa's distinctive solution of the Meno problem (e.g., Sosa, 2003, 2011). As is well-known, the Meno problem is a real challenge for reliabilist accounts of knowledge (e.g., Goldman, 1967, 1976; Kornblith, 2002). Why think that true beliefs formed through reliable competences are more valuable than not reliably formed true beliefs? Zagzebski (1996, 2003) vividly raised the problem by asking us to consider an analogy with two coffee machines, which on one occasion happen to deliver an equally excellent coffee. As it happens, one is an unreliable machine—it often produces bad coffee, just not on this occasion. The other is a reliable machine. Let's now focus on the value of the coffee they produced. If it is good, why care about the reliability of the coffee machine? *Mutatis mutandis*, if the output of a belief forming process is successful (that is, a true belief), why care if it was also produced reliably?

Sosa's virtue epistemology *is* a form of reliabilism, since it takes knowledge to be true beliefs that are due to reliable competences.[4] However, Sosa's approach arguably improves on reliabilism in that it offers new resources to answer the Meno problem. The distinctive normativity of actions is key. In the domain of actions, our practices of credit allocation are clearly valuing actions due to skills over actions due to luck. Compare a basketball player who sinks a basket from a difficult angle exhibiting all their adroitness to another who just gets lucky. The former deserves credit for their action, whereas the latter does not. The reason is that the action is due to their skills in the former case and to luck in the latter case. Thus, in our practice of credit allocation, whether or not an action is due to skills rather than to luck is an important consideration.

So, the action-theoretical angle of the virtue epistemological framework offers new resources for a solution to the Meno problem. According to this solution, knowledge is more valuable than true belief because performances that are due to skill are more valuable than those which are due to luck. The explanation for why knowledge is especially valuable—and at any rate more valuable than mere true belief—is predicated on the idea that only success due to skill is fully creditable to the agent. Hence, an important theoretical advantage of the virtue epistemological framework over reliabilist views of knowledge is predicated on whether **Creditability** holds.

[3] On another possible construal of creditability, the normative force of creditability is that of a *permission*: A success is creditable to an agent just in case the agent *can* be credited for it. On this construal, the creditability of a performance only licenses permissions of credit allocations, not requirements. Though Sosa is not explicit about it, I believe we should read Sosa as making the claim that aptness of a performance amounts to its agent's deserving credit for it—to the requirement construal of creditability.

[4] Indeed, the earliest statement of the virtue epistemological framework explicitly invokes the ideology of reliabilism (Sosa 1980: 23).

3 Creditability and skills

Let us look at **Creditability** more closely, then. According to **Creditability**, creditable performances are apt performances, where apt performances are performances that are correct in virtue of *the exercise* of skills. Here "in virtue of" stands for a (partial) grounding relation. Moreover, the "the exercise" clarification is important. Consider a scenario in which my being a rather skillful skier makes me very confident at skiing, so out of excessive confidence I undertake a very difficult ski jump I have never undertaken before, and end up succeeding at it due to mere luck. In this case, my success is due to my skills in *some sense*—in the sense that my skill led me to undertake the performance; but the success was nonetheless too lucky for me to deserve credit for it. For me to deserve credit for it, the success has to be in virtue of the *exercise* of the skill too: because the success is due to possessing the skill (since this possession made me overly confident) but is not due to its exercise, the relevant "in virtue of" relation does not obtain.

With this clarification, **Creditability** is very plausible for a variety of skilled performances. Consider, e.g., if a sprinter scores a time record due to excessive wind or due to performance-enhancing drugs, rather than to their skills as a sprinter. In these cases, the performance is not to be credited to them, and the time record should not be awarded, just as **Creditability** predicts.

Nonetheless, a *prima facie* alternative idea to **Creditability** is worth considering. Perhaps, what is fully creditable is not success due to the exercise of skills but rather success that is due to effort (e.g., Bradford, 2015). One might think, in other words, that it is the effort of the performer that matters for credit, rather than the exercise of the skill.

This idea is important. We typically do not credit people for qualities that they did not work for—e.g., while we might *praise* people for their fitness regardless of whether they have worked for it or not, we only really *credit* them for their fit bodies if they did something to achieve them. This feature of our practices at allocating credit seems to be justified by the idea that it matters to credit allocation whether the agent has *worked* for the qualities that they possess.

However, this idea is not necessarily at odds with **Creditability**, once the right understanding of skills is in place. Acquiring a skill *is* typically an effortful endeavor. Psychologists tell us that it requires *aimed* practice—practice with the goal of improving at the relevant task; when the relevant skill corresponds to an area of expertise, reaching expertise at that skill requires thousands of hours of practice (Ericsson, 2008). Similar considerations hold for *exercising* skills. Recent work on the philosophy of skills has dispelled the myth that exercising skills is effortless. Even when it appears effortless, skilled behavior requires the active and focused engagement by the agent—without which the flexibility of skilled performances and the sort of innovation skilled agents are capable of would not be possible (e.g., Christensen et al., 2016; Montero, 2016; Pavese & Beddor, 2023; Pavese forthcoming). Thus, when the appropriate conception of skill is in place, one can appreciate that the effort view of credit and the skill view of credit are not necessarily at odds, if both possession and exercise of skills is an effortful matter.

So the skill view will agree with the effort view on the central cases. One might object that the effort view and the skill view come apart in other important cases. Suppose one is working hard at learning some task for which one has not acquired the skill yet. Isn't the learning creditable to them? It might seem that only the effort view is compatible with this verdict. However, in learning, one is typically exercising *some skills*—those at learning the relevant skill. So the child is learning how to play tennis, and the learning is creditable to them since they are exercising some skills that they have—e.g., more basic embodied skills that are prerequisites for playing tennis.[5]

Does creditability of a successful performance require that this performance be *skillful*? Against this suggestion, in some cases, it seems that an action might be creditable to an agent even though the action is not really skillful. For example, one might make a mediocre risotto and yet be creditable for it ("it was not good but you should still get credit for making dinner"). Though this example might show that creditability does not require full-fledged skillfulness, it does not count against **Creditability** since it does not show that creditability does not require the exercise of skills. Here, the subtle grammar of "skilled" and "skillfulness" should be taken into account. The adjectives "skilled" and "skillful" are *gradable adjectives*, just like "tall," or "heavy." A gradable adjective *a* is satisfied by some behavior *b* in a context *c* only if *b* possesses a to a degree that exceeds *c*'s standard. For example, being tall requires not just having a height but having it to a degree above the contextual standards; accordingly, though every basketball player possesses some height but only some might count as tall. Similarly, being skilled requires not just having a skill but having it to a degree above the contextual standards. Thus, some behavior might exercise a skill (to some degree), without counting as skilled, or skillful, in a context. Thus, just as **Creditability** has it, creditability might require the exercise of skill but not necessarily to a degree above the threshold required for skillfulness.

4 A taxonomy of competences

Thus, **Creditability** is well-motivated when it comes to a variety of competent performances. Notice, however, that all the examples discussed so far intuitively motivating **Creditability** involve *particular* kinds of competences. Let us grant that all competences are reliable dispositions (Sosa, 2010: 465); even so, not every competence is alike. The sort of competences in the examples above are ones that are (i) acquired (typically with effort and practice) and are (ii) actively exercised when they are exercised at all.

[5] What if one has no skill at all yet when starting to learn a skill? This circumstance raises Aristotle's puzzle of learning by doing (*Nicomachean Ethics*, Book 2, Chapter 4, particularly 1105a18-1105a34). See Pavese (forthcoming, Chapter 12) for a solution. Some anticipations: at first, their beginner's success at A-ing will be only lucky or guided by an instructor, and so not creditable to them. But eventually, they will acquire some degree of relevant skill, and their successes at A-ing will be correspondingly creditable to that level.

Skills properly said are competences with this profile. Cognitive scientists, biologists, and philosophers alike even define "skills" as particular kind of competences necessarily involving learning (Dreyfus, 2004; Pavese, 2016; Ryle, 1949; Singleton, 1978; Willingham, 1998, Yarrow et al., 2009; Pavese forthcoming). As usually understood in the psychology and philosophy of skills, skills are acquired competences that require for the exercise the agent's active engagement: the archer intentionally acts when they shoot their arrow, and so does the basket player when they sink the basket. Agentive control—the sort of control characteristic of expert skilled behavior—has been theorized by cognitive scientists, as well as philosophers of psychology (Christensen et al., 2016; Fridland, 2017; Pavese, 2016; Pavese & Beddor, 2023; Wu, 2020). 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 initiate the task and to adjust the execution of the action as the performance unfolds.

However, not every competence (i.e., not every reliable disposition) has this profile. There are competences that are either not acquired or whose exercise does not require agentive control. Consider *instincts*. These are innate competences that are inherited from our lineage (e.g., Lorenz, 1957). They are typically *adaptations*—the result of natural selection—though they do not need to be. Their exercise does not typically result in agentively controlled performances. For example, in humans, blinking mechanisms in response to threats are instinctual. While these are very reliable competences, they characteristically manifest not in actions but in nonagentive behavior—e.g., blinking mechanisms are just reflexes (similarly for yawning mechanisms).

Yet other kinds of competences are acquired like skills but come apart from skills in other respects. Consider habits: like skills, they are also acquired. And their acquisition also constitutively requires practice. Precisely because they result from the agent's practice—which requires repeated actions—habits are more reflective of our agency than instincts are. Nonetheless, habitual behavior differs in a variety of important ways from skillful behavior and primarily in the kind of *control* that they exhibit. In psychology and neuroscience, habitual behavior is defined operationally as behavior that is *insensitive* to changes in the goals of a task (e.g., Haith & Krakauer, 2018:11). This operational definition is motivated by a variety of examples in which habitual behavior can come apart from the goal of the agent. For example, if the steering wheel is on the opposite side of the car, one may find yourself habitually reaching toward the door when trying to shift gears or pull the handbrake. Or consider the habitual nature of typing, which is similarly unmasked if one tries to type on a foreign keyboard, in which certain symbols might be mapped onto different keys (Hardwick et al., 2017); or cases of slips, as when, right after moving to a new place, you find yourself driving to the old one.[6]

These examples illustrate that habitual behavior might manifest control in some sense—e.g., in the sense that it can adjust to the unfolding to the circumstances to further the goal of the agent, when the goal of the agent is *fixed*; nonetheless, it

[6] See Amaya (2013) for an extended discussion of slips.

Table 1 A taxonomy of competences

	learned	Actively exercised
Skills	✓	✓
(Mere) habits	✓	✗
(Mere) instincts	✗	✗

fails to manifest control in another important sense—i.e., that of being sensitive to changes in the agent's goals.[7]

This diminished control of mere habits manifests itself also in the manner of their elicitation. For a vivid example of *merely* habitual behavior, consider the neurological deficit of *ideo-motor apraxia*—a deficit that affects complex skilled movements (e.g., Pavese, 2021; Sathian et al., 2011). What is distinctive to this deficit is that it makes a difference to whether the patient can perform a motor task if the task is environmentally cued or if the agent attempts to initiate the task in absence of an environmental trigger. An apractic patient might not be able to perform tasks such as *making the sign of the cross* when asked to do so but might perform the sign with no problem when entering a church; or they may not be able to pick up the phone when asked to do so but might be able to perform the action automatically when the phone rings.

Ideo-motor apraxia is interesting because subjects with this condition have lost the sort of control that is distinctive of skills, as they cannot even control the initiation of the task. When elicited in response to environmental triggers (the presence of the church, the ringing of the phone), their behavior is not even agentive—it is nothing more than a reflex. Nonetheless, these subjects have retained the (mere) habit, as they remain reliably capable of automatic responses that are elicited by the environment—as the cross signing case shows—even when these responses conflict with their intentions (Table 1). This discussion gives us the following preliminary taxonomy of competences:

This taxonomy is not meant to be exhaustive: as we will see in a bit (Table 3), there might be competences that do not neatly fit within it; also, with it, I do not mean to suggest that skills, habits, and instincts are sharply separated. Skills *do* involve habits—for example, one can only be skilled at tennis if one has habituated to certain motor routines, which have become reflexes; moreover, lots of so-called habitual behavior manifests *some* agentive control: when we go on the autopilot, we retain control over initiating and stopping the task, and we can adjust our behavior to new obstacles (Arpaly, 2000; Kalis & Ometto, 2021).[8] Even instincts and skills

[7] This is not to deny that habitual behavior *can* be agentive. We develop habits in order to reach our goals and manifestations of these habits can thus be seen as being directed towards the achievement of those goals (Kalis and Ometto 2019). In these cases, it just seems wrong to deny that the relevant habitual behavior is agentive and intentional.

[8] Elsewhere I argue that this sort of habitual behavior is agentively controlled to a larger extent than *merely* habitual behavior (Pavese forthcoming) on the ground that it involves the exercise of skills, not just habits.

are not sharply separated: some instincts can give rise to corresponding skills, in the sense that some innate characteristics can provide *the scaffolding* for acquiring a corresponding skill. For example, a variety of studies suggest that infants (between three and twenty four months) have a predisposition for rhythmic movement in response to music and other metrically regular sounds and infants as young as five months exhibit some tempo flexibility (e.g., Ilary, 2015; Zentner and Earola, 2010). These studies suggest that babies have a distinctive *musical instinct*—i.e., an instinct to respond to music and metrically regular sound with bodily movements, which might facilitate the acquisition of musical skills.

This threefold distinction between skills, mere habits, and mere instincts is therefore compatible with these platitudes—i.e., that skills and habits are instinctual to some extent, that skills are habitual to some extent, and that lots of habitual behavior is somehow controlled. All the present taxonomy is meant to clarify is that, while skills are not sharply separated from instincts and habits, nonetheless skills are not identical to *mere* instincts (since skills are acquired and their exercises are agentive), nor are they identical to *mere* habits (since skills require control over initiating the task, whereas *mere* habits, exemplified by the aforementioned apractic behavior, do not).

5 Credit and creditability

The point of the last section's taxonomy of competences is that certain differences among competences—in whether or not they are acquired, and as to whether or not they are actively exercised—might matter for credit.

Consider thinking of aptness as success in virtue of mere instincts and understanding aptness in **Creditability** in this way. The resulting claim that apt performance is creditable loses in plausibility. As Dretske (1993: 204) vividly puts it, we do not get to be credited for the sort of behavior enabled by innate mechanisms such as, e.g., blinking in response to threats; in such cases, credit goes to Mother Nature, not to us:

> So despite the fact that I blink and such behavior is a useful thing to do, despite the fact that I also think it a useful thing to do, the behavior is not a manifestation of intelligence. It is not a manifestation of intelligence because the behavior, though in conformity with thought, is not explained by thought. It is not governed by thought. I do not do it because of what I think. **We can, if we like, thank evolution for designing automatic reflexes that prevent injury in such cases, but we cannot thank ourselves. If this is to be counted as an intelligent arrangement, Mother Nature, not me, gets credit for the intelligence. To credit me with intelligence in these situations is a mistaken assignment of credit. It would be like crediting a sprinkling system with intelligence when it puts out the fire that would otherwise destroy it** (Dretske 1993: 204; my bold).

For more examples illustrating Dretske's point, consider the beautiful wild-song of zebra finches, which the experts deem to be instinctual (e.g., Fehér et al.,

2009). In its aesthetic properties, their song is comparable to that of the most skilled opera singer. But while their musical performance at the opera is creditable to the opera singer, a virtually indistinguishable performance in the forest is not creditable to the *individual* bird. Or consider a lion hunting a prey: this instinctual behavior is not something the individual lion should be held responsible for, since it is not something creditable to them rather than to their instinctual nature.

Holding that instinctual behavior of this sort is creditable to the individual is a bit like claiming that innate physical qualities such as whether one has a fit or an handsome constitution ought to be credited to them, even though one has not absolutely done anything to possess them. But as we have seen in §2, this is not how credit allocation works. Agents do not get credit for their height or for their innate handsomeness.

Examples such as these suggest that it is *learned* competences—the sort of competences individuals have to work for—that matter for credit. This is to say that **Creditability** is especially plausible when understood as **Creditability_1** rather than as **Creditability_2**:

> **Creditability_1** A performance is fully creditable to their agent just in case it is successful in virtue of the exercise of an acquired competence.
> **Creditability_2** A performance is fully creditable to their agent just in case it is successful in virtue of the exercise of a competence, whether acquired or not.

This is the first respect along which a more fine-grained taxonomy of competences might matter for the normative status of aptness. There is a second respect too. We can distinguish two dimensions to creditability. One dimension is *historical*: as we have seen, part of the reason why skilled performances are creditable to their agents is that the agents have *acquired* the relevant skills. We should get credit for our skilled performance because it is us who have acquired the skills and in so doing we have enabled our body to respond appropriately to the circumstances (Table 2).

So, historical creditability is a component of full creditability. It cannot be all there is to creditability, however. Suppose I and you have spent the same amount of time and effort learning how to play the piano. On the concert day, you play skillfully, while I underperform. If historical creditability were all there is to the creditability of a performance, we should get the same credit for our performances. After all, we both have exercised competences that we have acquired and we might imagine that we are generally even equally skilled. But intuitively, the one who played better that night should get *more credit,* just by virtue of having better exercised their

Table 2 Dimensions of creditability

Dimensions of creditability	Their nature
Historical creditability	Deserving credit for acquiring the competence
Exercise-creditability	Deserving credit for exercising the competence

skill. The historical model of creditability alone cannot vindicate these quite intuitive patterns of judgments.

What I think this shows is that its historical profile cannot fully capture the entire credit profile of a skilled action. Recall the ideo-apraxia patient, who retains the mere habit but who lost the skill: they cannot actively exercise their competences and their behavior is completely elicited by the environment. To be sure, some credit still goes to them. After all, they have acquired the competence which they can now exhibit, albeit now only in some circumstances. But in an important sense their performance is not fully creditable to them, since it no longer exhibits agency. In this sense, it is not creditable *for the exercise of the skill*—it is not "exercise-creditable." Similarly for slips: if I arrive at my old address out of a slip, when I intended to go at the new address instead, but the old address was actually where I was supposed to meet my friend but I had forgotten, I do not deserve full credit for having gotten to the right meeting place.

Thus, it seems that full creditability involves both historical creditability and exercise-creditability. If so, then **Creditability_3** better captures full creditability:

> **Creditability_3:** A performance is fully creditable to their agent just in case it is successful in virtue of the *active* exercise of an *acquired* competence.

6 On the possibility of instinctual and habitual knowledge

We have now made some progress towards understanding why aptness—success in virtue of skills—should go with creditability. Success that is due to competences that are acquired with effort and practice and whose exercise requires the active engagement by agents is the sort of success for which agents deserve full credit. By contrast, success that is due entirely to competences, such as instincts, that we have inherited from our lineage *and* that are not agentive is not obviously credited to the agents. With Dretske's words, mother nature gets the credit for that, not us. Finally, success that is due entirely to competences, such as mere habits, that are acquired with effort and practice but whose exercise does not require the active engagement of the agents, might only deserve some, not full, credit.

In this section, I want to explore the question whether knowledge by mere habits or by mere instincts is at all possible. If it was, the discussion in the last section might suggest that this sort of knowledge would not be creditable in the same way or to the same extent to which apt (skilled) knowledge is creditable.

Start with the possibility of merely habitual knowledge. Could one acquire knowledge by mere habit—completely automatically in response to environmental conditions, in the way an ideo-motor apraxia patient behaves by automatically reacting to environmental conditions? It seems so: consider a believer who lost the ability to initiate inquiry and to control the formation of their beliefs as a result of their competences but who retains the ability to form beliefs automatically in response to environmental conditions. Such believer might automatically form beliefs in response, e.g., to perceptual stimuli but fail to form beliefs by active reasoning and inquiry, as they retain no control over initiating the belief forming process. Just like

the ideo-motor apraxia agent, this believer would not exercise their agency when they do form their beliefs. As such they do not deserve exercise-credit. But as the discussion of the normativity of habits suggests, they might nonetheless deserve *some* credit—i.e., historical credit, as the agent has taken time and effort to develop the habit, and deserves credit for as much. (Indeed, certain domain-specific forms of perceptions, such as speech perception, might work as mere habits, see O'Callaghan, 2015:475; Stokes & Nanay, 2020:315; Pavese forthcoming.)

Thus, the possibility of habitual knowledge still fits the general virtual epistemological idea that knowledge is always creditable—albeit perhaps not always *fully* creditable—since knowledge out of habit is always at least creditable in the historical sense. So the possibility of habitual knowledge raises the theoretical possibility of distinguishing between kinds of knowledge on the basis of the degree of creditability that they afford. By contrast, the possibility of instinctual knowledge might raise a more substantive challenge for the general idea that knowledge is *fully* creditable.

As a case study, consider *social learning*—learning from others. In evolutionary psychology, it is widely thought that social learning is adaptive. If social learning were adaptive, then it would be plausibly instinctual in our species. The typical argument for the adaptivity of social learning in evolutionary psychology goes in two steps. The first is that learning—whether social or not—is more adaptive than non-learning. The second step is that social learning of a certain sort is more adaptive than merely individual learning.

Why should we think that learning is more adaptive than non-learning? It is helpful to compare learning processes to the process of inheriting traits due to the so-called "Baldwin effect." The Baldwin effect happens when a trait occurring in an organism as a result of its interaction with the environment becomes gradually assimilated into its developmental genetic or epigenetic repertoire. Only in particular situations is this sort of genetic assimilation more adaptive than having to learn things anew—those situations in which the relevant feature of the environment is fixed over many generations (e.g., Boyd & Richerson, 1994). In other words, if the environment does not change for a sufficiently long period of time, it "pays" to assimilate over having to learn anew. If the environment changes at medium tempo, however, then it pays to learn at each generation. So on the assumption that the environment does not change too slowly, learning is more adaptive than the Baldwin effect.

What about the adaptivity of social learning versus individual learning? For long, it seemed natural to think that social learning can be less costly when the task to be learned is very complex, and if the environment does not change too quickly. In these cases, social learning pays off more than individual learning—learning from others might expedite the process of acquiring information about the environment.

Rogers's (1988) paradox insinuated a doubt in this natural thought. He asked to consider the evolution of culture in a hypothetical species, the snerdwump, which inhabits a rapidly changing environment. Some snerdwumps cope with a varying environment by individual learning—they obtain, at some cost, information about the environment, and then exhibit the behavior most appropriate to that environment. Other snerdwumps employ social learning; they simply adopt the behavior of

some other individual— i.e., their "cultural parent". Roger's paradox is the observation that, when the social learners co-exist with the individual learners, social learning is very adaptive, but this situation changes as the individual learners become fewer in number and the society is dominated by social learners. In this case, social learning becomes less adaptive than individual learning, since without the presence of individual learners, the sort of information that it provides can be out of date.

Roger's paradox has challenged the thought that social learning is generally adaptive. However, subsequent discussion has assuaged these doubts about the adaptivity of social learning. As Enquist, Eriksson, and Ghirland (2007) argue, it is paramount to distinguish between *social learners* and *critical social learners*. A critical social learner is somebody who starts by socially learning a solution and then critically evaluates whether this seems to be an OK solution; if it is not OK, individual learning is tried. While social learning might not be more adaptive than individual learning in general, a population of critical social learners does have more fitness than a population of individual learners.

This discussion suggests that *critical* social learning at least is adaptive. If so, critical social learning might be instinctual for us, and if so, as critical social learners, we might be capable of merely instinctual knowledge. In other words, as we start learning from others, we might at first acquire knowledge merely by virtue of our instinct to critically learn from others. Indeed, anthropologists (e.g., Van Leeuwen et al., 2014) have offered wealth of evidence that children rely on critical social learning more than apes do, suggesting that an inborn reliance on critical social learning might even be distinctive of our species.

In response, one might contend that it is not clear that in the social domain, something can be *both* knowledge and fully instinctual. Moreover, something's being part of our lineage inheritance is not by itself evidence that it is "merely instinctual." For instance, something could be "maturationally natural" (McCauley, 2011) in virtue of inherited factors but be a habit or a skill in that domain. Classically innate capacities might be only the foundation on which we gain habits or skills.[9]

This is certainly true: it might be that for agents like us, innate capacities for social learning only provide the foundation for social habits and social skills. On this hypothesis, we cannot acquire knowledge from others without having first developed social habits and social skills out of our innate social capacities. Only then, will we be capable of acquiring knowledge from others.

Granted, the foregoing considerations fall short of establishing thay merely instinctual knowledge is actual. However, my argument does not turn on the actuality of merely instinctual knowledge but on its mere *possibility*. Could not we at least imagine agents that could learn from others out of their mere instincts? Could not agents like these be at least *possible*? A negative answer would be surprising. After all, we *do* yawn only out of instinct—no skill is required for yawning; we behave instinctually also when we blink in response to threat; merely instinctual behavior is in these cases both actual and possible. By parity of reasoning, it would be surprising if instinctual social knowledge—out of mere social and nonsocial instincts—were not even possible.

[9] I am grateful to a referee for this line of objection.

To make the possibility of merely instinctual knowledge a bit more concrete, consider a circumstance in which the social environment does not change very fast and where the practices of learning from others remain unaltered for many generations. In a similar, perhaps counterfactual, situation, Baldwin effect would have it that a variety of social skills would be assimilated as social instincts. Such assimilated social skills, now fully instinctual, could be used to form socially induced true beliefs by agents who find themselves in this new generation—true beliefs that could very much have the functional profile of knowledge, that is, the sort of state that agents rely upon for acting, reasoning, etc. Imagine that for these agents, general perception and the faculty of language are also purely instinctual, so its exercise is not agentive. Could not these agents acquire knowledge by exercising their social instincts? How can the possibility of merely instinctual knowledge be ruled out from the armchair?

It is the mere *possibility*, not the actuality, of merely instinctual knowledge that already raises a problem for the creditability profile of knowledge. Even if merely possible and not actual, such instinctual knowledge would not fit the profile of creditability underwritten by **Creditability_3**. To the extent to which we are after an analysis of knowledge—one that concerns believers in general, not just believers that have evolved in the contingent way we have—the mere possibility of instinctual knowledge is worth taking seriously.

7 A reassessment of Sosa's virtue epistemological framework

The main goal of this article has been exploring the relation between creditability and aptness. The positive suggestion so far is that we might make progress understanding the relation between creditability and aptness by inquiring more generally about how different kinds of competences and their exercise might underwrite allocation of credit. Whether or not a competence is acquired and whether or not a competence is actively exercised—as well as to the extent to which it is actively exercised—might matter for the credit that the agent deserves for the exercise of that competence.

It is now time to assess Sosa's framework in the light of this more fine grained taxonomy of competences. As observed at the outset, Sosa does not distinguish between competences on the basis of whether they are acquired or innate. However, famously, Sosa draws the distinction between "animal knowledge" and "reflective knowledge." Sosa often offers some basic kinds of perceptual knowledge as examples of animal knowledge. One might wonder how his distinction between animal and reflective knowledge would map into the current taxonomy.

Sosa's understanding of the distinction between animal and reflective knowledge seems to have to do with whether the competence exercised is "high-order" or not as well as with the extent to which it is agentive. For example, Sosa (2007: 24, 32) (see also Sosa 2021:55) tells us that animal knowledge is apt belief, reflective knowledge is *aptly apt belief*—it is apt belief that one aptly believes. Relatedly, Sosa (2021:54) characterizes the distinction in terms of whether the knowledge involves the exercise

of **subconscious and subcredal** competences versus the exercise of **rational reflective** competences:

> At a *merely* "animal" level, the eye-exam subject does epistemically well by getting it repeatedly right through a **competence that operates with impressive subconscious, functional reliability**. **So, a sort of sub-credal competence** there provides a certain level of knowledge (in an ostensible guess aided by zero confidence). Yet this falls below the level of human judgment constituted by **our rationally reflective faculties**. Such *mere* animal knowledge thus falls below the reflective knowledge familiar to normal humans over a vast domain wherein we judge, share information, and coordinate action, through our remarkable ability to communicate (my bold).[10]

So understood, 'animal knowledge' is closer to 'instinctual knowledge' and to 'habitual knowledge' than it is to 'skilled knowledge', since it manifests less agency than skilled knowledge. However, it does not map neatly into either. It differs from instinctual knowledge in that Sosa thinks of animal knowledge as still in some sense a manifestation of agency (see especially Sosa & Cahallan, 2020): even though the result of subcredal and subconscious competencies, animal knowledge is still described as agentive.[11] Animal knowledge might also differ from habitual knowledge in that animal knowledge, at least the basic sort that is acquired through basic sense perception, is not obviously due to an acquired competence. After all, the senses and some general (that is, not domain-specific) aspects of perception are plausibly parts of our basic innate endowment. So some animal knowledge is formed through competences that are plausibly at least partly innate natural faculties (such as sense perception or memory) but which Sosa seems nonetheless to take to be at least minimally agentive.

Reflective knowledge, on the other hand, is arguably closer to skilled knowledge, in that in exercising the agent's rational and conscious faculties, it clearly exercises agency. Nonetheless, some differences are worth emphasizing. One might presumably actively exercise a skill without exercising sophisticated "reflective" faculties. For example, one might actively exercise one's skill as an investigator, even when one does not actively reflect on the reliability of one's skill. Or one might actively exercise one's skill at gymnastics while being mostly oblivious as to one's reliability. Likewise for skilled knowledge: it is a sort of knowledge formed skillfully but not necessarily reflectively. Moreover, if skilled knowledge corresponded to reflective knowledge, we would have to say that nonhuman animals that lack our sophisticated reflective faculties could not attain knowledge out of skill. But this does not seem right: many nonhuman animals might well have skills, understood as learned competences that require their active engagement for their exercise.

So it seems that reflective knowledge does not correspond neatly to skilled knowledge; we might say instead that it is a *particular* kind of skilled knowledge, one that involves the exercise of sophisticated reflective epistemic faculties. This discussion

[10] See also Sosa (2001:94), who distinguishes animal knowledge from reflective knowledge in that only the latter requires the believer to take an "epistemic perspective on the reliability of one's sources".

[11] For the idea that a faculty such as perception is agentive, see Noë (2004) and Siegel (2020).

Table 3 A more fine-grained taxonomy of competences

	Learned	Actively exercised	Actively and reflectively exercised
Higher order skills and/or reflective faculties	✓	✓	✓
Skills	✓	✓	✗
(Mere) habits	✓	✗	✗
Animal (e.g., perceptual) faculties	✗	✓	✗
(Mere) instincts	✗	✗	✗

gives us the following even more fine-grained taxonomy of competences (as for the first taxonomy in Table 1, the distinction between animal faculties, habits, instincts, skills, reflective skills is not meant to be exhaustive nor sharp, as in some cases competences might be *partly* habitual, partly instinctual, etc.) (Table 3):

Now that we have a more fine-grained taxonomy of competences and kinds of knowledge, we are in a better position to reassess Sosa's framework. There are two dimensions along which our discussion matters for this purpose.

First off, recall that Sosa attaches *full* creditability to aptness (§2; e.g., Sosa, 2021: 64). However, in the light of our more fine-grained taxonomy, notice that full creditability might be too high a standard to expect of *all* kinds of knowledge. One might have to allow for different levels of credit to go along with different kinds of knowledge: Habitual knowledge comes with *some* credit (historical credit) and so does perhaps animal faculty knowledge (exercise-credit), while skilled knowledge and reflective knowledge only come with full credit—*both* historical and exercise-credit.

This suggestion is a small, hopefully welcome, refinement of Sosa's framework. A bigger theoretical challenge might come from the possibility of *merely instinctual knowledge*—a sort of knowledge that results out of competences that are neither actively exercised nor acquired. Recall the analogy with our blinking or yawning mechanisms. While in every other case of knowledge—habitual knowledge, animal faculty knowledge, skilled, and reflective knowledge—it makes sense to talk of some level of credit, by contrast, I suggested that the normative profile of instinctual knowledge as deserving credit is presently up for grabs. In this respect, I see two theoretical options available to virtue epistemologists, both of which are worthy of investigation.

A first theoretical option is to allow for the possibility of instinctual knowledge but at the same time to do work towards motivating a notion of creditability that is more permissive in that it recognizes some, however minimal, credit to instinctual knowledge. Accordingly, one might argue that despite its non-agentivity and its being unlearned, agents deserve some credit for their instinctual knowledge, e.g., via mere attributability.

While this option is theoretically available, it comes with some costs. As we have seen, we do not credit agents for their hair color or for their innate fitness. Nor do

we credit agents for their instinctual yawning or blinking. Why would things be so substantially different for knowledge formed out of mere instincts? This knowledge would be merely instinctual behavior, and so presumably only creditable to the extent blinking and yawning are creditable.

The second problem is that the resulting watered-down notion of *deserving credit* might not be fit for the theoretical bill. Recall that **Creditability** plays a central role in Sosa's distinctive solution to the Meno problem (§2). Now, the issue arises as to whether instinctual knowledge is more valuable than true belief—whether instinctual knowledge "passes" the Meno problem test. If, when it comes to assessing the coffee produced by the reliable coffee machine, the fact of how reliably it was produced seems irrelevant, why think, instead, that when it comes to assess the value of instinctual knowledge over and above that of true belief, the reliability of the instincts the knowledge is due to is relevant? Of course, in some sense, the coffee is more attributable to the reliable coffee machine than to the unreliable one. And yet this consideration does not seem to matter when assessing the value of the coffee. *Mutatis mutandis* for merely instinctual knowledge. The problem is principled: since instinctual knowledge is not even agentive, one can no longer respond to the Meno problem that arises for this kind of knowledge by importing the distinctive normativity of actions (§2).

The second theoretical option seems more promising. This is to insist that, as Dretske (1993: 204) has it, credit does not attach to merely instinctual behavior. If one adopts this option, then one might contend that animal knowledge maps into habitual, skilled, and animal faculty knowledge but not into instinctual knowledge. This route would amount to rejecting the possibility of merely instinctual knowledge. This move need not be entirely ad hoc. In order to motivate it from a theoretical point of view, one might start from taking knowledge to be whatever epistemic state that passes the Meno problem test; and from there go on to argue that, for it to pass the Meno problem test, knowledge cannot be merely instinctual. It seems to me a route worth exploring, one that has a lot to recommend it.

Acknowledgements I am grateful to Adam Carter, Jonathan Dixon, Chris Kelp, Mike Hicks, Daniele Sgaravatti, Mona Simion,Gus Turyn, the audience at the Urbino Summer School in Epistemology 2024, as well as an anonymous reviewer of Philosophical Studies, for discussion on topics related to this paper.

Conflict of interest The author declares no conflict of interest.

References

Amaya, S. (2013). Slips. *Noûs, 47*(3), 559–576.

Arpaly, N. (2000). On acting rationally against one's best judgment. *Ethics, 110*(3), 488–513.

Beddor, B., & Pavese, C. (2020). Modal virtue epistemology. *Philosophy and Phenomenological Research, 101*(1), 61–79.

Boyd, R., & Richerson, P. J. (1994). Why does culture increase human adaptability? *Ethology and Sociobiology, 16*, 125–143.

Bradford, G. (2015). *Achievement*. Oxford University Press.

Christensen, W., Sutton, J., & McIlwain, D. (2016). Cognition in skilled action: Meshed control and the varieties of skill experience. *Mind & Language, 31*(1), 37–66.

Dreyfus, S. E. (2004). The five-stage model of adult skill acquisition. *Bulletin of Science, Technology & Society, 24*(3), 177–181.

Enquist, M., Eriksson, K., & Ghirlanda, S. (2007). Critical social learning: A solution to Rogers's paradox of nonadaptive culture. *American Anthropologist, 109*(4), 727–734.
Ericsson, A. K. (2008). Deliberate practice and acquisition of expert performance: A general overview. *Academic Emergency Medicine, 15*(11), 988–994.
Fehér, O., Wang, H., Saar, S., Mitra, P. P., & Tchernichovski, O. (2009). De novo establishment of wild-type song culture in the zebra finch. *Nature, 459*(7246), 564–568.
Fridland, E. (2017). Skill and motor control: intelligence all the way down. *Philosophical Studies, 174*(6), 1539–1560.
Goldman, A. (1967). A Causal Theory of Knowing. *The Journal of Philosophy, 64*(12), 357–372.
Goldman, A. (1976). Discrimination and perceptual knowledge. *The Journal of Philosophy, 73*(20), 771–791.
Greco, J. (2000). Two kinds of intellectual virtue. *Philosophy and Phenomenological Research, 60*, 179–184.
Greco, J., DePaul, M., & Zagzebski, L. (2003). Knowledge as credit for true belief. In M. DePaul & L. Zagzebski (Eds.), *Intellectual virtue: Perspectives from ethics and epistemology* (pp. 111–134). Oxford University Press.
Green, A. (2016). *The social contexts of intellectual virtue: Knowledge as a team achievement*. Routledge.
Haith, A. M., & Krakauer, J. W. (2018). The multiple effects of practice: Skill, habit and reduced cognitive load. *Current Opinion in Behavioral Sciences, 20*, 196–201.
Ilari, B. (2015). Rhythmic engagement with music in early childhood: A replication and extension. *Journal of Research in Music Education, 62*(4), 332–343.
Kalis, A., & Ometto, D. (2021). An Anscombean perspective on habitual action. *Topoi, 40*(3), 637–648.
Kornblith, H. (2002). *Knowledge and its Place in Nature*. Oxford University Press.
Lackey, J. (2009). Why we don't deserve credit for everything we know. *Synthese, 158*(2007), 345–361.
Lorenz, K. Z. (1957). The nature of instinct. In C. H. Schiller (Ed.), *Instinctive behavior the development of a modern concept* (pp. 129–175). International Universities Press.
McCauley, R. N. (2011). *Why religion is natural and science is not*. Oxford University Press.
Montero, B. (2016). *Thought in action: Expertise and the conscious mind*. Oxford University Press.
Nanay, B., & Stokes, D. (2020). Perceptual skills. In E. Fridland (Ed.), *Routledge handbook of the philosophy of skill and expertise* (pp. 314–323). Routledge.
Noë, A. (2004). *Action in perception*. MIT.
O'Callaghan, C. (2015). Speech perception. *Handbook of the philosophy of perception*, 475–494.
Pavese, C. (2016). *Skills in epistemology I and II*. Philosophy Compass.
Pavese, C. (2021). Practical concepts and productive reasoning. *Synthese, 199*(3), 7659–7688.
Pavese, C., & Beddor, B. (2023). Skills as knowledge. *Australasian Journal of Philosophy, 101*(3), 609–624.
Pavese, C. forthcoming. *The Practical Mind. Skills, Knowledge, and Intelligence*. Cambridge University Press.
Riggs, W. D. (2002). Reliability and the value of knowledge. *Philosophy and Phenomenological Research, 64*, 79–96.
Rogers, A. R. (1988). Does biology constrain culture? *American Anthropologist, 90*(4), 819–831.
Ryle, G. (1949). *The concept of mind*. Routledge.
Sathian, K., Buxbaum, L. J., Cohen, L. G., Krakauer, J. W., Lang, C. E., Corbetta, M., & Fitzpatrick, S. M. (2011). Neurological principles and rehabilitation of action disorders: common clinical deficits. *Neurorehabilitation and Neural Repair, 25*, 21S-32S.
Siegel, S. (2020). Skill and expertise in perception. In E. Fridland & C. Pavese (Eds.), *The routledge handbook of philosophy of skill and expertise* (pp. 306–313). Routledge.
Singleton, W. T. (1978). Laboratory studies of skill. In W. T. Singleton (Ed.), *The analysis of practical skills* (pp. 16–43). Springer Netherlands.
Sosa, E. (1980). The raft and the pyramid: Coherence versus foundations in the theory of knowledge. *Midwest Studies in Philosophy, 5*(1), 3–26.
Sosa, E. (2001). Human knowledge, animal and reflective. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition, 106*(3), 193–196.
Sosa, E. (2007). *A virtue epistemology: Apt belief and reflective knowledge* (Vol. 1). Oxford University Press.
Sosa, E. (2010). How competence matters in epistemology. *Philosophical Perspectives, 24*, 465–475.
Sosa, E. (2011). *Knowing full well* (p. 129). Berlin Studies in Knowledge Research.
Sosa, E. (2015). *Judgment and agency*. Oxford University Press.
Sosa, E. (2021). *Epistemic explanations: A theory of telic normativity, and what it explains*. Oxford University Press.
Sosa, E., & Callahan, L. F. (2020). Skill and knowledge. In E. Fridland & C. Pavese (Eds.), *The routledge handbook of philosophy of skill and expertise* (pp. 146–156). Routledge.

Van Leeuwen, E. J., Call, J., & Haun, D. B. (2014). Human children rely more on social information than chimpanzees do. *Biology Letters, 10*(11), 20140487.
Wayne, W. (2020). Automaticity, control, and attention in skill. In E. Fridland & C. Pavese (Eds.), *The routledge handbook of philosophy of skill and expertise* (pp. 207–218). Routledge.
Willingham, D. B. (1998). A neuropsychological theory of motor skill learning. *Psychological Review, 105*(3), 558.
Yarrow, K., Brown, P., & Krakauer, J. W. (2009). Inside the brain of an elite athlete: The neural processes that support high achievement in sports. *Nature Reviews Neuroscience, 10*(8), 585–596.
Zagzebski, L. (1996). *Virtues of the mind: An inquiry into the nature of virtue and the ethical foundations of knowledge*. Cambridge University Press.
Zagzebski, L. (2003). The search for the source of epistemic good. *Metaphilosophy, 34*(1–2), 12–28.
Zentner, M., & Eerola, T. (2010). Rhythmic engagement with music in infancy. *Proceedings of the National Academy of Sciences, 107*(13), 5768–5773.