Teleological Epistemology: Knowledge-how and Goal-directedness

Kye Palider

Abstract
Epistemology characteristically operates at an intellectual level, being primarily restricted to the cognitive states of human individuals and the factors that influence these states. This restriction may be undesirable for a more generalized epistemology that discusses non-individual or non-human systems that may be said to be intelligent, such as communities, animals, and complex systems whose parts may not be independently intelligent. Here, I show how such a generalized epistemology may be possible under a teleological, non-intellectual, and not conventionally mechanical view of knowledge. I establish a teleological schema of knowledge-how, where a system is understood to know how to X if it is capable of being goal-directed towards X. This schema is analogous to how philosophers of biology are increasingly describing organismal agency. After developing this view, I show how it handles many controversial cases of know-how in the literature. I then show how knowledge-that may be constructed out of this teleological know-how, showing that all knowledge may be understood teleologically, without the need to reference to individual human cognitive states. Thereby, I show a generalized approach to epistemology through teleology that can extend naturally to topics in social epistemology and non-human intelligence.

Keywords: knowledge-how, epistemology, intellectualism, teleology, philosophy of biology, goal-directedness, Wittgenstein

Introduction

Traditionally, epistemology operates at an intellectual level as a discipline, even going so far as to take epistemic, cognitive, and intellectual to be synonyms (Harrison 2013, 199–200). The objects of study of epistemology involve, exclusively, cognitive objects such as beliefs, desires, memories, doubts, certainties, intentions, etc., which are understood in an idealized and

perhaps opaque way. These cognitive objects are elements of a phenomenological model of our experience. Such a phenomenological model cuts our mind at agreed upon joints: certain thoughts or experiences are labeled as beliefs, some as desires, and even potential thoughts may be given labels, such as memories. At such an intellectual level, epistemology provides an ontology which is to be evaluated by how well it agrees with our individual experiences and how well it adheres to our linguistic usage of these terms. Notably, such an epistemology, by virtue of being framed at the intellectual level, is an epistemology of the individual for the individual, where any social–epistemological constructs must supervene over individual experiences, and any sociocultural factors are only relevant insofar as they influence the cognitive state of the individual.

This intellectual level is typically contrasted with a more natural level. This natural level is equated with science, empiricism, or even mechanism (for historical reasons), and is thought of as non-intellectual. Natural objects behave mechanically, being governed by laws of nature or certain regularities. They do not feel or experience, they do not think of possibilities, and they do not weigh options or make choices; they are prima facie non-intellectual. A tension arises if we come to accept that our very cognition supervenes, in some sense, over natural mechanisms. How can the intellectual be constructed from seemingly non-intellectual parts? This question and related questions are for the most part a topic in the philosophy of mind; however, I would like to focus on a question that (also) belongs to epistemology: can knowledge, as the central topic of epistemology, be understood non-intellectually? This question may come off as incoherent or oxymoronic at first as knowledge, if anything, must be intellectual, but what I mean more specifically is if we can have an epistemology that does not operate from an intellectual–phenomenological–individual level. And, if knowledge can be understood non-intellectually, must it be understood naturally along the lines of neuroscience, the cognitive sciences, or neurophysiology?

One might say that neuroscience has already made strides in this regard. Around a century ago, after the consolidation of much experimental work, we came to understand our nervous system as a network of cells called neurons that live in an electrochemical soup and fire electrical signals to one another according to a certain calculus. Working off this neuron doctrine, we have come to know many processes of nervous systems, such as how a visual cortex may discern between colours. However, as far as neuroscience has come, we still have very little to say about the composition of a belief, a desire, an intention, or effectively any of the cognitive concepts of epistemology. One might respond that it is only a matter of time before we discover the

composition of beliefs, desires, and all the other concepts of epistemology. Yet, there is a sense in which this response is misguided. Even if we were to fully map the human nervous system, and understand the composition and dynamics of neurons and other relevant structures, including their immediate environment, we would still face difficulty in identifying what process or part of this 'complete' picture of the human intellect is a belief. Could we circle some neurons and say that they are, or somehow contain, your belief that $1 + 1 = 2$?

One might object that although such a project is daunting, we have a method of arriving at conclusions of such a kind. We could run experiments that stimulate the brain and see which stimulation results in the subject reporting that they have experienced a belief that $1 + 1 = 2$. To this, I present many concerns. Some are standard scientific concerns: perhaps beliefs are realized in many heterogeneous ways, requiring the neuroscientist to identify the common structure between the neural states that correspond to the report of belief; and perhaps they are realized differently between individuals, requiring the neuroscientist to find the common structure between individuals. Other concerns, the neuroscientist cannot resolve. Perhaps the subjects are working off different understandings of a belief, so how would we tell they share a (sufficiently) common understanding or, more importantly, the right understanding of what it means to experience the belief that $1 + 1 = 2$? Whether they share a common understanding may be tackled by the linguist who assesses their usage of the term 'belief' and its semantics, so it is still resolvable. However, what makes it the 'right' understanding is not an empirical question and is in the domain of the philosopher. Even if philosophers were to agree to a correct meaning of the term 'belief' and identify it with an intellectual kind, what guarantees that there is a corresponding natural kind?

There may be no corresponding natural kind. Although a term such as 'gold' may correspond to a natural kind (an element with a certain atomic number), perhaps the same cannot be said about belief or other cognitive concepts. There is good reason to believe this is the case as our understanding of these cognitive concepts seems to permit them being realized in heterogeneous ways by other agents, such as an alien (cf. the Martian intuition in Sprevak 2009), perhaps bearing no nervous system like ours. If these cognitive concepts are not natural kinds, then they may be more akin to functional kinds (Fuller 2022) or what I will call a teleological kind. For instance, a wheel can be realized in many ways, made out of stone, out of wood, out of rubber and metal, etc., and it is still a wheel so long as it is able to fulfill its function of revolving around an axle and generating traction

against a surface. A wheel is not a natural kind.[1] Furthermore, it is not, as a concept, phenomenological (although the sight or thought of a wheel may be). The wheel operates in between, so to speak, these two levels as a teleological kind. It is my goal to argue that the same can be said of knowledge.

I approach this challenge through the proxy of knowledge-how (know-how). Know-how concerns knowing how to carry out certain tasks or achieve certain ends. It is to be contrasted with knowledge-that (know-that) which typically refers to having a cognitive state that is related in a certain way to a proposition, which means that you know that proposition (to be true). The central debate concerning know-how is whether or not it reduces to a form of know-that—whether or not know-how is cognitive/intellectual or not, or whether it is propositional or dispositional (Bengson & Moffett 2012b, 163). Intellectualists (such as Bengson & Moffett 2012b; Stanley & Williamson 2001; Stanley 2011; Schroeder 2012) argue that know-how can be reduced to knowing certain propositions, which leads to a form of propositionalism about knowledge if we assume that know-how and know-that exhaust all forms of knowledge. Anti-intellectualists (such as Ditter 2016; Ryle 2009; Noe 2005; Hornsby 2016; Fridland 2013) argue that know-how is separate from know-that as it is not propositional, but dispositional. This disposition is typically framed as follows: a subject knows how to φ if and only if they have the ability to φ intentionally (Cath 2019, 489). Consider knowing how to ride a bike. The anti-intellectualist would say that one knows how to ride a bike if they have the ability to ride a bike intentionally, with some success and with the qualification that the circumstances are appropriate or normal (Bengson & Moffett 2012a, 170). The intellectualist would deny at least one direction of this if and only if, and claim that knowing how to ride a bike is instead knowing ways of riding a bike, or propositions that answer questions such as "What is a way I can ride a bike?". The intellectualist is so-called as she operates at the intellectual level, with cognitive states such as beliefs, while the anti-intellectualist is so-called because she operates outside the intellectual level, with dispositions or abilities. The intellectualism debate is then at the frontier between an epistemology at the intellectual level and one outside of it, making it the decisive battle that determines whether or not knowledge can be viewed non-intellectually. As a result, succeeding in showing that know-how does not operate at the intellectual level is sufficient for showing that some branch of epistemology need not be intellectual, and further succeeding in showing that

[1] That is not to say it cannot be studied by the natural sciences. The natural sciences may be concerned with how to construct wheels, the properties of wheels, the mechanics of wheels, etc. However, the concept of a wheel would only be developed as a non-empirical endeavour, by the inner 'philosopher' of the natural scientist.

know-that need not be interpreted along an intellectual, phenomenological, or propositional dimension would be sufficient for showing that epistemology, as a whole, can be understood differently, along some other dimension. One such other dimension, I argue, is teleological.

My argument is as follows. First, I identify the underlying teleology present in existing accounts of know-how (Section 1). I then present a teleological understanding of know-how, akin to teleological concepts in (the philosophy of) biology (Section 2). The philosophy of biology is especially relevant for this case as it presents a view of agency as teleological that does not subsume anything inherently cognitive or intellectual about agency, nor does it subsume any spooky non-natural forces in its description of organismal agency. After developing the teleological understanding of know-how, I show how it evaluates traditional and controversial cases of know-how (Section 3). Lastly, I show how one may construct know-that out of a teleological know-how (Section 4). This takes inspiration from Wittgenstein's notion of rule-following and statements regarding knowledge, which has been recognized as having a significant relation to know-how (Harrison 2013). If all these steps succeed, then I will have shown that there is a non-intellectual, teleological form of know-how that know-that reduces to, and as such all forms of knowledge can be viewed from a non-intellectual perspective, requiring no grounding in phenomenological, individual, or propositional factors. Epistemology, so teleologized, takes the emphasis off the individual, human subject and opens the door to a more general epistemology of goal-directed systems that need not be individual, nor necessarily restricted to human cognitive states and their influences.

Teleology in Knowledge-how

Teleology is not at all foreign to the intellectual level. In fact, it is plausible that our first-hand experiences of intending to do something are wherefrom we derive our notion of teleology in the first place. Know-how is no stranger to teleology as accounts of know-how often formulate it explicitly with respect to intentions. Here, I briefly sketch the pervasiveness of teleology in the literature on know-how.

As noted earlier, anti-intellectualists typically adopt a dispositional view to know-how. Such a dispositional view is usually put similarly to the following: a subject knows how to φ if and only if they have the ability to φ intentionally (Cath 2019, 489). Although there exist many variations of the anti-intellectualist formulation of know-how, often adding further qualifications to at least one direction of the previous statement, the intentional aspect of it is integral. One would be hard-pressed, for instance, to say that someone knows

how to catch a butterfly if a butterfly flew into their net while they were distracted, but were unable to catch a butterfly when they were actively trying to. This ability to intentionally φ also typically requires some degree of counterfactual success (Bengson & Moffett 2012b, 170). Even if the butterfly flew into the wanna-be bug catcher's net while they were actively trying to catch the butterfly, if such an event was an improbable accident and in all similar cases they would have failed to catch the butterfly, one would once again be hard-pressed to say that the wanna-be bug catcher knows how to catch butterflies. Most anti-intellectualists, as such, require the ability to intentionally φ in order to know how to φ and some degree of counterfactual success, or a condition that shows the intention to φ was in some sense a cause of the result of φ-ing.

Anti-intellectualists then clearly have a teleological understanding of know-how. For an agent to intend to φ, it is their goal or telos to φ; they act towards achieving an end state where they φ. This teleological component of know-how is essential; without it, all successful instances of φ-ing appear to be accidental. Nuances between anti-intellectualist views typically appear with respect to the form of counterfactual success required, and these shall be surveyed shortly, after presenting the teleology present in intellectualist views.

Intellectualists, beginning with Stanley and Williamson (2001), took knowing how to φ to be reducible to knowing the propositional answers to questions of the form "How do you φ?" or "What is a way of φ-ing?". Know-how is then a relation between a cognitive state, a belief, and a proposition, representing a valid way ω to φ. Note that even if one is an intellectualist, one may still accept the biconditional of the anti-intellectualist; the difference would be that the ability to intentionally φ is not something non-cognitive, but holds in virtue of purely cognitive factors, such as having certain beliefs. However, most intellectualists would deny the biconditional on the grounds that one can know how to φ without being able to φ at all. Take, for instance, a pianist who has had their hands amputated (Stanley & Williamson 2001, 416). The proposed intuition is that the pianist still knows how to play the piano despite being unable to play. Propositionally speaking, nothing has changed with respect to the pianist's cognitive state on matters relating to ways of playing the piano; the only thing that has changed is her ability to play the piano. The intellectualist may then say that know-how has nothing to do with abilities and all to do with cognitive states, thereby restricting know-how to the cognitive level.

Despite the intellectualist denying know-how from entering the non-intellectual level of dispositions, they would still claim that know-how is action-guiding. For instance, Cath (2019, 491) notes that a common theme in

both the intellectualist and anti-intellectualist literature is that know-how involves norms regulating actions and abilities. In their non-propositional intellectualist view, Bengson and Moffett (2012b, 176) state that "An attractive suggestion is that knowledge how is potentially action-guiding in the sense that it is a state that can guide successful, intentional action". The force here is that the pianist may have a cognitive state that can, if they were to have hands or be in typical conditions, lead to successful, intentional playing of the piano. In this sense, even if know-how reduces to a cognitive state, it is still a cognitive state that can be action-guiding towards a specific goal (performing the action one knows how to do). It is in this sense that know-how is inherently goal-directed or teleological even for the intellectualist.

For know-how to be goal-directed or teleological, it must also be resilient with respect to obstacles that appear in pursuing that goal; it cannot just regularly achieve its ends by accident. This level of resilience may introduce a gradable notion of know-how, which may be based on the persistence and plasticity with which the agent pursues their goals (Sommerhoff 1950 and Nagel 1979 originally proposed the terms 'persistence' and 'plasticity'). For instance, many people may know how to bake a cake while following a recipe, using specific equipment, and having access to specific ingredients. However, the best bakers know how to bake a cake according to many recipes, with different equipment, and with varied ingredients. The best bakers' know-how is action-guiding in a more robust sense than the know-how of amateur bakers as it leads to a successful end state (a good cake) in more distinct scenarios. This teleological feature of know-how is present for anti-intellectualists in general through some form of counterfactual success, where the broader the degree of counterfactual success (perhaps the more nearby possible worlds in which the intention to φ successfully leads to φ-ing), and for intellectualists who demand that a cognitive state regulate actions in pursuit of the goal to φ.

This teleological feature is also readily apparent in Gilbert Ryle's work that first introduced the modern conception of know-how in 1949:

> To be intelligent is not merely to satisfy criteria, but to apply them; to regulate one's actions and not merely to be well-regulated. A person's performance is described as careful or skillful, if in his operations he is ready to detect and correct lapses, to repeat and improve upon successes, to profit from the examples of others and so forth. He applies criteria in performing critically, that is, in trying to get things right. (Ryle 2009, 17)

Greater know-how can better overcome obstacles, and it is learned by repeating and improving upon successes, as well as by learning from example. This adds a dimension to know-how that is not purely cognitive or intellectual, but causal or teleological.[2] Elzinga (2021) takes this self-regulated view of know-how and extends it more naturally to animals that may or may not have the same cognitive or representational resources as us, such as rats and honeybees. Since these action-guiding or teleological features of know-how are not at the human intellectual level, they may be used to model the behaviour of animals. Rats, for instance, can be taken to employ map-based strategies when navigating their environment in pursuit of food, or even in labyrinths set up in a laboratory (Elzinga 2021, 1746). Even though it remains an open question as to whether rats have propositional knowledge, it is natural to say that they have some form of know-how in the form of being able to persistently and plasticly pursue goals (such as procuring food) in the face of many potential obstacles. Elzinga (2021, 1755) then presents his view of know-how as follows: "To self-regulate is to try to get things right. Trying to get things right in this context may be analyzed as goal-directed behavior controlled by a representation of the goal state, which activates the process of making trials and terminates with accomplishing the goal (at least when things go well)".

I take Elzinga's formulation of know-how to be quite general here. It seems to be applicable to many agents that are not necessarily human, who may or may not have attitudes towards propositions. However, Elzinga, in this goal-directed or teleological formulation of know-how, still requires that that behaviour is controlled by a representation of the goal state. For humans, many of our representations are captured through beliefs in propositions. For rats, there may be good evidence that they think in terms of maps. For honeybees, the story may be different, and it may be different for all sorts of organisms. The requirement that there be a representation of the goal state still carries with it the weight of the intellectual level; it is still an opaque concept that refers more to the purported phenomenology of the agent as opposed to a more natural, causal, or teleological understanding of know-how. In the following section, I motivate an extension of Elzinga's account of know-how that drops the condition of representing the goal state. For this, I take inspiration from the growing literature on biological agency, which does not need to make use of any cognitive states to explain agency. This then leads to an entirely teleological and non-intellectual view of know-how.

[2] Note that the causal and teleological need not conflict here. It is causal in the sense that know-how may be embedded in a belief–desire model in order to predict future action. It is teleological in the sense that the agent may be modeled as pursuing a goal, perhaps a desire but not necessarily, and learning to pursue that goal better. This will be expanded upon in the next section.

Teleology and a Teleological Schema of Knowledge-how

The teleology common in talk of intentional action does not lend itself well to a natural interpretation. The agent is taken as an external, intangible decision maker, whose choices manifest themselves in the more tangible actions they take. Such a view of teleology may inadvertently lead it to being labeled a dualist, vitalist, or simply supernatural enterprise, as has historically been the case, especially since the Renaissance (Gambarotto 2019, 152). The shift to a mechanistic understanding of the universe, as seen in Descartes, was in part a rejection of the teleological approach in the Aristotelian paradigm, and the acceptance of a laws-of-nature-governed worldview. The natural sciences now consider goals and purposes unacceptable (Uller 2023, 2). Instead, natural phenomena are to be explained by causes, forces, and mechanisms. The human being came to be understood as a system of mechanical parts, and any intentions it might have are simply the result of all its parts' non-intelligent, law-governed activity. This mechanistic understanding of the world and of ourselves became pervasive, and is now an underlying assumption at the backbone of the natural sciences—one that denies teleology a role because of its supposed unnaturalness.

However, although teleological theories such as vitalism evidently conflicted with a certain physical understanding of the world (e.g., one with conservation of energy), teleology has historically existed symbiotically with mechanical views of nature. Aristotle's philosophy that emphasized teleology as a form of natural purposiveness defined the Western world's study of nature, or natural philosophy, for over a millennium (Gambarotto 2019, 154). In a way, for Aristotle, denying a teleological understanding of nature is simply a deprived view of nature, as all one is doing is cutting out one of the causes, final causes, in favour of a view dominated by efficient causes. The dominant theories of physics could be said to have made do without final causes, but other sciences could be said to have suffered. It is now becoming increasingly apparent that a mechanistic view of biology is not an ideal perspective, or at least one that misses out on key features of complex biological systems (Walsh & Rupik 2023). Perhaps teleology, in general, offers many advantages for a study of systems that are complex, chaotic, not understood in a reductive fashion, and which feature processes that are opaque to our current understanding, such as economics (viewing people as utility-maximizing rational agents), biology (which will be expanded upon below) and the study of human cognition (as it relates to know-how or knowledge in general).

A mechanistic understanding of biology generally favours what comes to have been called the gene's-eye view (Ågren 2021), where agency and the actions of individual organisms are determined from below by their genes and

from above through population dynamics. Organisms behave the way they do because they are so propelled by forces from the genes below and the population dynamics above. If a certain population of organisms comes to fracture a bone, one would be inclined under the gene's eye view to say that they have a predisposition to fracture this bone, or a gene for fracturing this bone. Many take this view to be unilluminating, as it ignores the agency of the organism or reduces the agency of an organism to the downstream consequences of its genes. It is likewise unhelpful to view organisms as genes reacting to their immediate environment and packing all that the gene or the organism does during its lifetime into a single quantity of 'fitness', as is done in Fisherian population genetics. This is to be contrasted with an organism-first view, or agential view (for a history and extensive bibliography, see Baedke & Fábregas-Tejeda 2023).

Biologists and philosophers of biology are starting to recognize that organisms may be best viewed as agents that pursue goals (Desmond & Huneman 2020, 33). Instead of saying that a beaver carries a gene for dam-building, it may be easier to say that the beaver is an agent that pursues goals such as having a shelter, protecting itself from predators, storing food, etc., and that in pursuit of these goals, it makes use of its causal repertoire (what it can do; in this case, primarily make use of its teeth) to build a dam. Organisms are not just genes reacting to their environment or samples of a statistical distribution; they are agents that act of their own accord and create their environment. Organisms, be they human beings, beavers, or even single-cellular organisms, pursue goals that are typically beneficial to their survival.

This has led philosophers of biology to formulate what it means to be an agent in a way equally applicable to human beings and bacteria. Although there are several views on exactly how to spell out this organismal agency, it is generally agreed that an agent is a goal-directed system that pursues goals of its own initiation (Walsh & Rupik 2023, 8). Organisms are agents that can also be taken to be intelligent in a sense of the term: "[Agency] consists in the system's capacity to transduce, configure, and respond to the conditions it encounters. Crucially, agential systems are capable of maintaining functional stability in response to conditions that would otherwise compromise their viability" (Sultan et al. 2022, 5). Their functional stability is their ability to pursue their goals persistently and with plasticity in the face of obstacles.

Crucially, note that this form of natural agency does not take it to be an intellectual phenomenon. An organism can be a goal-directed system without having cognitive states per se. The organism need not have the ability to form intentions, have desires, or be sentient whatsoever; it is understood teleologically as a goal-directed system without reference to the teleology of

intentions at the intellectual level. It is also not a characteristically mechanical view. It does not explain the function of the organism through its constituent parts, down through the cells, to the proteins, and eventually to the atomic level; it explains the function of the constituent parts as regulated by the whole system's goals. Organismic agency is a natural, ecological property rather than an intentional property (Desmond & Huneman 2020, 40). One can also add that it is a teleological property that is non-intentional and operates outside the intellectual–phenomenological level.

The similarities between natural agency and know-how are immense. Going back to Elzinga's self-regulation view, a system may be said to know how to φ if it can self-regulate itself based on a representation of the goal to φ and initiate a process that terminates in it φ-ing. Just like natural agency, Elzinga requires a goal-directed system that initiates processes towards the goal. There are, however, two main differences between natural agency and Elzinga's know-how. First, natural agency typically requires some form of organizational closure (Jaeger 2024, 6). A system is organizationally closed if all essential parts of the system mutually depend on each other for their continued existence. In this sense, although a human body may be organizationally closed, the human nervous system is not, as it depends on further parts to maintain its continued existence, that is, it would not continue to exist unless it is embedded in a human body that is essential for its continued existence.

I take the first difference to be one unique to the aims of biology as a discipline. Organizational closure is useful in identifying what makes an organism a continued living thing, what is part of an organism, and what is not essential to an organism. However, it does not seem to be relevant in a more general sense to the ability of a system to pursue its goal; a system may pursue a goal as a component that depends on a relatively stable broader system, but nevertheless may be responsible for the broader system's pursuit of that goal. For instance, the lungs are not organizationally closed but may robustly pursue the goal of supplying the bloodstream with oxygen and carrying out carbon dioxide, which is an integral part of the human body's goal of maintaining its continued existence. In this sense, organizational closure may be of interest for biologists interested in what makes an organism an agent throughout time, but it need not be for other disciplines such as epistemology where goal-directedness itself may be the object of interest.

The second main difference between Elzinga's view and natural agency is that Elzinga requires a representation of the goal state, which natural agency does not. The purpose of a biological system is not defined through any representation the system may bear as a cognitive feature, but is defined

loosely as a state of the system (Desmond & Huneman 2020, 40). In such a way, even single-cellular organisms such as bacteria have a purpose even though they do not have a nervous system and even if we do not know exactly how that purpose is manifested through the underlying features of bacteria. The behaviour of bacteria is simply best modeled as though they were pursuing goals, robustly at that, even if we do not understand the mechanisms responsible for them behaving as such.[3] As such, Elzinga's requirement of a representational state is more stringent than what is required for natural agency.

What if one were to drop this requirement for know-how and try to formulate a purely natural and non-intellectual definition of know-how? This question motivates the remainder of the paper. First, let us try to formulate what it means to know how to φ if one drops the condition that it requires a representation of the goal state:

Teleological Know-how. A system knows how to φ if the following three conditions hold:

1. The system is capable of being goal-directed towards φ-ing.
2. The system's being goal-directed towards φ-ing initiates a process that, with some persistence and plasticity, results in it φ-ing.
3. The first two conditions hold in certain contextually relevant environmental conditions.

The first condition encapsulates the notion that the system can pursue φ. For a person, this might simply mean that they can have the intention to φ, but it more generally means that some state of the system causes the system to behave as though it were trying to achieve the goal of φ-ing. The second condition requires that the system's goal-directedness is explanatory of it φ-ing and that its φ-ing is not a mere accident. Furthermore, the degree of persistence and plasticity is in some sense analogous to the level of skill or expertise with which the system knows how to φ. How to quantify the level of persistence and plasticity may be context-dependent. Returning to the example of the expert baker, their know-how can initiate a process that leads to the baking of a cake in many more distinct scenarios than the know-how of an amateur baker and while facing many different obstacles. In some contexts, perhaps being flexible with the ingredients you use is more important than the equipment, while in others, the opposite may hold true. The last condition is

[3] In this sense, a teleological understanding abstracts away from the micro-level mechanisms of the system and describes its behaviour more generally, as a system with a purpose.

also context-dependent. For instance, one would not expect anyone to be able to bake with a lack of atmosphere, or perhaps if they are sick or inebriated. As such, these conditions are somewhat flexible and can be interpreted as a schema of know-how rather than a strict formulation.

Although all of these conditions are quite vague, so too may any linguistic usage of know-how be vague. At times, we may require a high success rate in φ-ing. Other times, we may appeal to typical environmental conditions that are non-actual. For instance, an astronaut could be said to know how to skip rope even though they cannot skip rope while aboard a space station. We may, furthermore, deny know-how if the agent can only φ in particular ways, requiring some plasticity on behalf of the agent. There's more than one way to skin a cat. Hence, there are at least these three dimensions along which the contextual standards of know-how vary.

The implications of such a teleological view on know-how are perhaps best seen through the lens of many examples. The following section explores teleological know-how through several classic examples, as well as some new examples that broaden the know-how literature. Topics in the epistemology of non-individual non-human systems is especially relevant here. These topics may seem to be of no immediate interest, but they are central in, for instance, the social epistemology of communities who make use of instruments. In fact, many of these topics will have biological analogues that shed light on what exactly makes a system intelligent without reference to any cognitive states it may have. This further identifies the common ground between the intelligence of communities, of biological organisms, and of the human brain.

Increasingly Radical forms of Knowledge-how

Teleological know-how may be advanced as a position in two ways. First, it can be a claim that the use of the term ‘know-how’ linguistically bears its semantics. I pursue this aim by offering a flexible reading of know-how, which hopefully restores its semantics in an intuitive manner. However, the second way may be to take it as a concept that has a life of its own, with its own utility in describing knowledge and intellect. Wittgenstein put it well:

> Some will say that my talk about the concept of knowledge is irrelevant, since this concept as understood by philosophers, while indeed it does not agree with the concept as it is used in everyday speech, still is an important and interesting one, created by a kind of sublimation from the ordinary, rather uninteresting one. But the philosophical concept was derived from the ordinary one through all sorts of misunderstandings,

> and it strengthens these misunderstandings. It is in no way interesting, except as a warning. (Wittgenstein 1980, 289)

In either case, let's begin with some examples.

> **Pianist two ways**. (1) A skilled pianist loses their hands in a tragic accident. (2) A musical theorist who never had hands has studied how individuals play the piano to a deep level, understanding all the hand positions and movements required to play many pieces. She receives new (sci-fi) prosthetics that respond to her thoughts in a predictable fashion: whenever she thinks about playing certain keys with a certain hand position, the prosthetic plays those keys with that hand position. She successfully plays the piano with the prosthetics.

The accepted intuition in the literature is that the pianist knows how to play in case (1). Teleological know-how, under a standard reading, however, results in the handless pianist not knowing how to play the piano. Although they can intend to play the piano, and hence become a system with the goal of playing the piano, they cannot initiate a process that reliably results in them playing the piano. If know-how is taken to be a cognitive state, then it is clear that the pianist's cognitive state with respect to their piano playing is unaffected by the loss of their hands. No change in cognitive state, no change in know-how. However, presuming such a statement would be to beg the question on an intellectual view of know-how, and perhaps our judgments should run against these intuitions.

We may say that, under this reading, the pianist only knew how to play the piano with certain tools (their hands), but now lacking the relevant tools, they are no longer able to play the piano. This is akin to an amateur swimmer only being able to swim with a pool noodle, and now deprived of that pool noodle, one would hardly say that they can swim. Compare this to case (2), where an individual studies the hand movements required to play the piano without actually ever playing it. They are not naturally said to play the piano, but once they receive a prosthetic that is somehow responsive to their conceptions of the hand movements required to play the piano, they are able to play the piano. In (2), the individual received tools that they could use with their existing cognitive state in order to play the piano, and were, as a result, able to play the piano. They know how to, under this reading, play the piano.

One might disagree with the above interpretation and say that the pianist case fails to reconstruct an accurate representation of the semantics of know-how. However, here is an alternative reading. Suppose instead of focusing on

the individuals as entire human systems, we instead restrict ourselves to their nervous systems. We further stipulate that the contextually relevant environmental conditions are those with a 'normal' (in this case, two-handed) human body. With these stipulations, we see that the pianist's nervous system in (1) would reliably be able to play the piano if it so desired, if it were 'hooked up' to a typical human body with hands. The pianist in (2), however, would not know how to play the piano under such a reading because their nervous system has not learned the dexterity required to move typical human hands, even though it could move the prosthetic.

I take this specification of the teleological know-how schema to be a common one, especially for intellectualists. Under this specification, what does the knowing is the human nervous system, or perhaps the brain, and it is what we evaluate with respect to being goal-directed. The environment is specified as what can or is typically attached to that nervous system. In the case of the pianist, we may assume human hands. In other cases, we may assume some contextually relevant body. Take, for instance, the following example.

> **Ski Instructor** (Stanley & Williamson 2001, 416). A ski instructor regularly teaches Olympic athletes how to do difficult stunts. The instructor cannot, themself, do the stunts as they are not sufficiently athletic; however, they were able to do the stunts when they were younger.

In this case, the ski instructor's nervous system, if attached to a sufficiently able body, such as that of its young self, would be able to do the stunts. So, we may take these environmental conditions to be the relevant ones. As such, we may simultaneously accept that the ski instructor knows how to do the stunt, but does not know how to do the stunt *with his current body*.

Another lesson can be learned from the second pianist case. It seems that there exists a tool, even if highly idealized, available for any system (such as a cognitive state) to achieve its goal. For instance, consider someone who has no idea how to play the piano, but they are wearing a suit that responds to their natural language queries and moves their body to meet their demands. They command it to play a piano concerto, and off it goes. In this sense, their cognitive state knows how to play the piano under the environmental conditions of being inside the suit. However, outside of the suit, they are unable to play the piano and as such cannot be said to know how to play the piano without the suit. This leads me to deny a common formulation where "knowledge how is potentially action-guiding in the sense that it is a state that can guide successful, intentional action" (Bengson & Moffett 2012b, 176). If

the 'can' is interpreted as mere possibility, then there could exist so-called smart tools that enable any so-called dumb agent to achieve their goals, such as the suit of the previous example. The can should hence be weakened to contextually relevant scenarios in order to reconstruct a more semantically accurate formulation of know-how, assuming this is the end Bengson and Moffett (2012b) are going for.

It should further be noted that these cases can also be interpreted through the lens of different systems. We may, for instance, take the individual and the suit to form a goal-directed system together. The heavy-lifting in terms of playing the piano would be done by the suit, but the intention and goal-directedness of the system would come from the individual. It might then be natural to say that, most of the time, when we say someone knows how to φ, we restrict the goal-directed system to either their body as a whole, or just their nervous system. Let us take a look at another similar case.

> **Neurological disorder salchow** (Bengson & Moffett 2007, 46). Irina is a figure skater. She believes that to perform a quintuple salchow she needs to take off from her back inside edge and land on her back outside edge. In reality, she needs to do the opposite: take off from her back outside edge and land on her back inside edge. However, she has a neurological disorder where she mixes up the back inside edge with the back outside edge. Her belief in conjunction with her neurological disorder reliably leads her to perform quintuple salchows.

The intended intuition of this example is that Irina does not know how to perform the quintuple salchow because she does not know a valid way of performing a quintuple salchow. So, even though she regularly performs quintuple salchows, she does not have the know how. If one wishes to resist this intuition, it would be quite easy to apply the teleological schema. Irina, as a whole human, or even whole brain, can be goal-directed towards performing a quintuple salchow, and when she is so goal-directed, she initiates a sequence of movements that result in her performing it. It is illuminating to see how one might reconstruct Bengson and Moffett's (2007) intuitions. There seems to be a privileging of those cognitive or neurological states that contain the beliefs of ways of performing a quintuple salchow as opposed to Irina's entire body, nervous system, or even brain. In one interpretation, one part of Irina's neurology, which may be said to 'contain' the belief regarding the inside and outside edge, uses the rest of Irina's neurology (with the neurological disorder) to successfully achieve its goal of performing a quintuple salchow. Although it is dubious if one can compartmentalize a nervous system as such, let us grant

it for now. This system, as a whole, then knows how to do a quintuple salchow. In another reading, the restricted neurology that contains the purportedly problematic belief, if attached to a normal remainder of a nervous system, would not reliably lead to the performance of a quintuple salchow. Hence, this latter reading reveals something that may be implicit in intellectualist views of know-how: know-how is restricted to the goal-directed system that carries the beliefs of the agent (and perhaps desires, or whatever is required to formulate intentions). This does seem to, again, beg the question in the intellectualist debate.

Here is an example that questions the first condition of teleological know-how:

> **Liar-turned-truthteller**. After a life-changing event, a skilled, experienced, and successful liar makes a complete 180 in their worldview and vows to never lie again, even if it were to mean one's own death or worse. She still remembers and deeply regrets all the lies she's said before, remembering the methods by which she constructed her lies. The only sense in which she can lie anymore is by accident (if one were to still call this lie as opposed to a mistake) or requires some element of forcing, for instance, someone attaching electrodes to certain parts of her mouth muscles or the corresponding nerves and forcing her to utter a lie. Otherwise, she has a will of iron and will never lie.

Does the liar-turned-truthteller know how to lie? In some sense, she does; she knows many ways of lying, she could explain the ways one might lie, and she could teach others to lie skillfully. However, in another sense, she cannot get herself to do it—she cannot intend to lie, i.e., make lying her goal. It is in this sense that one might say she does not know how to lie.

I take this example to pull apart two competing conceptions of knowing how to do something. On the one side, call it the intellectual side, the liar-turned-truthteller knows the answers to questions such as "How can you lie?" and "What are some ways one can lie?". However, on the flip side, what some would call the anti-intellectual side, the liar-turned-truthteller does not have the capacity to lie due to her will of iron, or, one might say more broadly, factors external to the propositional knowledge but relevant to action. With respect to a teleological account of know-how, she cannot orient herself as an agent such that her actions are directed towards the goal of lying. Although it is possible for her to lie, it would not be initiated by her as a process directed towards the goal of lying, but by someone or something else's extrinsic purposiveness. Even if she lived in a world where all forces conspired to make

her lie and she ended up lying on many occasions despite many obstacles being in place to prevent her from lying, for her to know how to lie, it is required that she can make it her intrinsic purpose to lie, and not that she becomes an accessory as an extrinsic purpose to lie, no matter how robustly she lies.

Now, let's move on to a non-human example.

> **Neutrophil**. A neutrophil is a type of white blood cell. It is able to identify bacteria and viruses, and then chase them down. Bacteria try to evade the neutrophil, but it is able to chase bacteria past other cells and can track bacteria even if the bacteria are 'hiding' behind other cells or around corners.

A neutrophil is a classic example of a natural agent according to philosophers of biology. It maintains organizational closure while pursuing goals of neutralizing bacteria as part of the larger immune system. It achieves its goals with reliability as it is able to chase bacteria and then consume them despite the best efforts of the bacteria (which are themselves agents) in attempting to evade the neutrophil. There are many videos (from a microscope) of neutrophils engaging in such chases, which I recommend watching—they are convincingly intelligent, appearing as though they are a child playing tag with the bacterium. Does the neutrophil know how to neutralize bacteria? Under the teleological view of know-how, they evidently do as they are goal-directed systems that can reliably and robustly achieve their goals in the contextually relevant environment of the human body. They know how despite not having a nervous system and only consisting of a single cell.

Overall, teleological know-how gives justice to the view that there are intelligent organisms beyond human beings and it clarifies what makes them intelligent. They are intelligent because they are capable of pursuing goals even in the face of obstacles. One can further begin to quantify their intelligence by discussing the kinds of obstacles they may overcome. Sometimes, the obstacles may be simple to overcome, and other times they may require the system to adapt. For instance, the immune system is constantly updating its 'directory' of bacteria for neutrophils to chase in response to certain bacteria intruding upon the body. This adaptation can be viewed as a form of learning without a nervous system.

Here is a pair of examples that share a similar structure:

> **Collective know-how**. (1) Scientists working at CERN are part of a large organizational structure involving experts in many different fields. Together, they perform experiments and present their results. (2)

> Different groups of neurons part of the human nervous system can be said to serve different functions. The visual cortex is able to discern between colours, various nerves traveling through the spine are able to sense heat at the extremities they reach, etc. Together, these groups of neurons form the nervous system of someone who knows how to do many things.

The first case comes from Knorr-Cetina (1999), where she discusses the high-energy physics community conducting experiments through particle accelerators that require the coordination of teams of physicists, engineers, and other staff, along with sophisticated instrumentation, in order to conduct experiments involving particle collisions. Knorr-Cetina claims that all the people part of CERN and the instruments form a system of distributed cognition. All the research divisions and sets of instruments mutually depend on one another in order for the system, as a whole, to be goal-directed towards conducting high-energy physics experiments successfully. Although some research groups may be goal-directed themselves, they would not succeed without all the other groups and the instruments. Even if we take the set of all people working at CERN, they would still fail to conduct their experiments without the instrumentation.[4] CERN as a whole consisting of people and instruments, and only as a whole, knows how to conduct the precise high-energy physics experiments that it does. It is a goal-directed system without necessarily having an easily interpretable set of cognitive features.

This evaluation is also in alignment with the broader literature on collective or joint know-how. Birch says that "joint know-how arises when each agent knows how to predict, monitor, and make failure-averting adjustments in response to the behaviour of the other agent, while actively enabling the other to make such adjustments" (Birch 2019, 3329). These conditions are precisely those that would make the joint group a goal-directed system that can robustly achieve its goal by making 'failure-averting adjustments'. Similarly, Palermos and Tollefsen (2018, 121) require that the contributing members of a system that knows how "collaboratively perform a cognitive task by interacting continuously and reciprocally with each other". This is again a form of mutual interdependence that is very common in views of natural agency. In fact, the parallels to biology run deep, as one can model cells as units that interact

[4] One might shift around what is part of the system and what is part of the environment. One could say the people know how to use the instrumentation in order to perform their experiments, but I do not take this to deny that the system of people and instruments as a whole knows how to conduct the experiments simpliciter, without the qualification that the instrumentation is part of the environment.

continuously and reciprocally with one another, and the unit cells (which are agents themselves) form a larger agent (the multi-cellular organism) collectively. Let us now turn to the second case, which is of this kind.

Case (2) here is owed to Professor Joseph Berkovitz in conversation. Much like how the individuals working at CERN and the instruments involved in the high-energy physics experiments perform their respective functions that help the system as a whole attain its goal, systems of neurons may also perform their respective functions that help the nervous system achieve its ends. Subcollections of neurons know how to perform subtasks of the nervous system's entire cognition, and groups of neurons mutually depend on one another to perform the task of overall cognition. In this way, the nervous system as a whole may be taken to be a goal-directed system that can pursue many goals by virtue of its simpler parts, which each individually know how to perform certain functions. Accordingly, human cognition can be viewed as the manifestation of non-cognitive, or at least less-cognitive, systems working together in a goal-directed fashion. This argument can be further extended to the components of individual neurons, eventually bottoming out with some form of atoms. Thereby, the goal-directedness of a nervous system and of neurons themselves may supervene over the arrangement of entirely non-cognitive parts.

It should be noted, however, that there is a sense in which atoms or fundamental particles know how. If fundamental particles are governed by laws of nature, then it may be said that they know how to follow the rule specified by the law of nature. They are de facto goal-directed towards obeying a law of nature as they cannot do otherwise. They will with absolute persistence and plasticity obey the law of nature. And, they will do so in all environmental conditions. In this way, they know better than anything else possibly could how to abide by the laws of nature. Ascribing know-how to fundamental particles in this way may be quite unintuitive. One could try to add further conditions to teleological know-how, such as the ability to pursue distinct goals, and define know-how through those additional conditions, but they seem to be ad hoc solutions, if they work at all. The know-how of fundamental particles may be a bullet one simply needs to bite. This might be a case in point where know-how and intelligence diverge. Perhaps intelligence is marked by the capacity to pursue many distinct goals or to pursue goals of a certain kind (complex ones, for instance), or a certain way of modulating a repertoire in pursuit of a goal. In any case, the schema of teleological know-how proposed implies that law-governed activity constitutes know-how, but this might only be unintuitive if we conflate know-how with more general intelligence.

Over the course of these examples, teleological know-how has been expounded. It is a schema that can, under certain specifications, reconstruct the intuitions of many cases in the literature. However, it often is most naturally read as going against those intuitions. Despite its initial counterintuitiveness, through multiple readings of the schema, it reveals the underlying assumptions behind widely held intuitions across the literature. Often times, intellectualists isolate the cognitive states responsible for belief and intention from all other aspects of an agent, and say that know-how must reside in those cognitive states supplemented with a typical human body and conditions. What I hope to have shown is that we do not need to assume know-how to exist in the corners of these cognitive states; instead, know-how is ubiquitous in living organisms and even in many compound systems. This runs counter how many authors view living organisms and automata, such as the following:

> On one hand, such knowledge how seems to be practical—unlike mere knowledge that, which can be possessed even by incompetent, impractical "fools". On the other hand, knowledge how seems to be a genuinely cognitive, even if not a ratiocinative or discursive, achievement—unlike mere abilities or dispositions to behavior, which can be enjoyed even by mindless entities or automata, such as simple machines and plants. (Bengson & Moffett 2012b, 161)

Although I agree that know-how is more substantive than simply assenting to a belief, I disagree that 'mere' abilities or dispositions are unintelligent. (Plants are agents and can robustly pursue various goals that are conducive to their broader goal of survival). A teleological understanding of know-how broadens the scope of epistemology to, perhaps more accurately, capture what it means to be intelligent outside individual human intelligence. The dispositional states of 'mindless' entities such as automata and plants may, in fact, be intelligent in the sense that they are dispositions that lead those entities to persistently and plasticly pursue goals.

The Knowledge-how of Knowledge-that

The account of know-how argued for herein may seem anti-intellectualist in all respects. It argues that know-how is not something that is best understood at an intellectual level and that know-how has more to do with being a goal-directed system than it does with beliefs in any propositions. However, there is one aspect of the intellectualism debate that I have not touched upon: is know-how distinct from know-that? The traditional intellectualist argument is that know-how is simply a form of know-that, or that it reduces to it. Here,

based off insights from Wittgenstein, I argue that, if anything, the reduction runs opposite: know-that is a form of know-how, that is, propositional knowledge is reducible to procedural knowledge.

There is already a sense that I have covered in which know-that may be seen to reduce to know-how. Recall the case of groups of neurons serving functions as part of the broader nervous system, which is capable of being goal-directed towards many different goals. If beliefs, desires, and intentions are cognitive states that supervene over our neurology, then know-that, representing a subset of those beliefs that stand in a particular relation to the rest of the world, also supervenes over our neurology. One can then make several claims, such as the nervous system knowing how to believe or know a proposition, or the nervous system knowing how to act according to a belief. It might be the case that, further, some subset of the nervous system is the one that handles a specific belief, and it can be said to know how to do these things. Although this line of reasoning sounds plausible, I find it unilluminating.

We are simply reasoning about neurological states while assuming that the belief is reducible or composed of that neurological state. However, there may be reason to believe that a belief cannot be taken to be a natural kind whatsoever. There is the experience of a belief, which many would take to be a non-natural, intellectual kind, but there is also the idea that beliefs could manifest for agents with different, non-neurological architectures (Sprevak 2009). If this is right, then it suggests that a belief can be interpreted as a functional or teleological kind, similar to a wheel. The intellectual and phenomenological reading of a belief may not be argued to be reducible, owing to anti-physicalist concerns (such as those in Chalmers 2004). However, we have yet to assess whether know-that, if it is taken to be a teleological kind, is reducible to another teleological kind, namely, know-how.

How might such a reduction look like? I take it to be more akin to a functional decomposition. Returning to the wheel example, note that a wheel can be decomposed into several essential parts that each serve their respective function. Something that functions as an axle, something that rotates around the axle, and something that generates traction against a surface seem to be necessary components of a wheel. We would not be able to have a wheel without these components. A teleological reduction, in this sense, can be interpreted as a functional decomposition. The broader goal is reduced to subgoals of the component parts.

In order to assess whether such a reducibility is at all feasible, we must first start with a teleological understanding of belief. I present two, what I hope seem plausible, options for a teleological understanding of belief. The first is a pragmatist understanding of belief, as originating from Frank Ramsey's work.

The second is a linguistic or representational understanding of belief. In both cases, I show that it is know-how all the way up to know-that.

In *Facts and Propositions* (Ramsey 1927), Ramsey denies the canonical viewpoint that what makes something true is some relation it has to a fact, thereby abandoning a correspondence theory of truth. He characterizes correspondence theories of truth as supposing that a certain relation holds between beliefs, judgments, or assertions—what he calls the mental factors—and the objective factors which could be propositions, facts, or objects. Beliefs, or the mental factors more generally, for Ramsey and the larger school of pragmatist thought, are understood causally. Having certain beliefs, in the context of certain desires, is simply to be disposed to act in a certain way. What makes a belief or a system of beliefs true is for it to be maximally expedient at meeting desires, in a general sense. The true belief system is the most expedient belief system for all agents of all desires or at least equivalent to other belief systems in expedience, but never exceeded. What realizes the expedience of a belief is understood through the objective factors; the world must be organized in some particular way in order for the beliefs to be causally efficacious at maximizing the desires of agents with those beliefs. Later, in *Truth and Probability* (Ramsey 2016), Ramsey expands upon the causal role of beliefs through constructing the now-ubiquitous belief–desire model. As such, Ramsey is committed to a 'thin' conception of truth, where he does not specify a 'thick' notion of correspondence between the mental and objective factors, but still requires that they be aligned in some fashion.

Belief so understood emphasizes that beliefs are the causal precursors to (some) actions. What actions beliefs lead to is determined through the belief–desire model, which can be quantified through decision theory, a field Ramsey inaugurated. The beliefs and desires, as cognitive states, serve as mechanisms or efficient causes for the agent's actions and intentions. A belief–desire model provides a causal–mechanical account of an agent's actions.

Such a causal–mechanical understanding of beliefs, desires, action, and, as a consequence, knowledge is actually not so different from the teleological one advanced herein; they are like two sides of the same coin. Both describe the same causal processes, just in two different ways. The pragmatist view cites the mechanisms involved in beliefs and desires as the causes of action. The teleological view focuses on the goal-directedness of a system and explains actions as pursuits of the goal. The desires serve as final causes while the beliefs initiate efficient causes towards fulfilling the desires. Both views could be causally equivalent in the sense that they agree on the same counterfactuals. "Had the agent's goals been different, or had their causal repertoire been different, they would have acted differently" may be translated to and from

"had the agents desires been different, or had their beliefs been different, they would have acted differently" without loss of meaning, where the resulting actions are identical. In this way, these two conceptions of belief are two perspectives, or models of belief, that may agree on underlying predictions, counterfactuals, or orders of events and possible events, but just disagree on the best terms to explain that underlying structure, be it teleological or mechanical. As such, the teleological and pragmatist theories of action are empirically equivalent (van Fraassen 1980; Quine 1975; Laudan & Leplin 1991), even though they differ in their ontological commitments and heuristic value. One theory may be simpler or have better pragmatics (in the linguistics sense of the term) for a given situation, which may mean that it is preferable for that situation, but neither theory is better in all respects. For instance, it seems that some disciplines prefer a mechanical understanding of how things happen such as physics, while other disciplines, such as those to do with human actions and biology, may prefer a teleological understanding of how things happen.

Now, let us turn to the second linguistic conception of belief. Here, a belief can be taken to be a statement of a proposition. More generally, this proposition may be said to be some representation of the world (or possible worlds), such as a map-based understanding of the environment, or something like a set of possible worlds in which the proposition is true. It is important to note that this conception of belief's 'linguistic' nature only requires that the belief be formulated along some language, loosely defined. It may be some form of language of thought or mentalese that varies, for instance, between organisms that bear beliefs. Under such an understanding of belief, what it means to believe something is to assent to that representation of the world. In a typical human case, this would be to assent to a statement that you claim to believe in.

The intellectualist may further extend this notion of assent to a modal one. If the following conditional holds, then the agent may be said to believe in the referenced proposition: if the agent were to entertain a proposition (or representation of the world), then they would assent to it. This can be seen in typical accounts of the implicitness or intuitiveness of know-how. Knowing how to ride a bike is complicated, and it may be hard to believe that the intellectualists require that you know, explicitly, ways of riding a bike including the minutiae of your movements required to keep the bike balanced and moving. Instead, the beliefs are implicit, in the sense that they would be assented to if entertained. This explains the frequent automaticity of know-how as opposed to know-that, which may not take away from its intentional and intelligent nature (Fridland 2017).

Such a conception of belief takes beliefs to be communicable, expressible, or at least capable of being assented to. I work off Wittgenstein to show that such a formulation of belief leads to know-that that must reduce to know-how. Consider the following: "One has already to know (or to be able to do) something before one can ask what something is called. But what does one have to know? […] We may say: it only makes sense to ask what something is called if he already knows how to make use of the name" (Wittgenstein 1968, 30–31). Although Wittgenstein is focused on the name and its meaning, this remark holds equally for a representation and its meaning. Before one can communicate, express, or assent to a representation, one must understand it or be capable of acting as though it bears such a representation. In order to genuinely assent to a proposition, or know that, one must first understand the proposition. Understanding the proposition involves knowing what it implies, how it fits in with other beliefs, how it is justified, and what rules govern its application—before one knows that a proposition is true, one must know how to act as though that proposition is true. Without such an understanding based on know-how, claiming to know that is just an empty utterance, devoid of any relevant action-guiding content.

There is a further parallel with another remark by Wittgenstein. Wittgenstein here is focused on ways of escaping an infinite regress of justifying an inference, owed to concerns raised by Carroll (1995), but it equally applies to the formation of a belief or cognitive state more generally:

> Giving grounds, however, justifying the evidence, comes to an end; but the end is not certain propositions' [Sätze] striking us immediately as true, i.e. it is not a kind of seeing on our part; it is our acting, which lies at the bottom of the language-game. (Wittgenstein 1979, 204)

Here is my interpretation of this line of reasoning with respect to belief: the formation of a belief is a process that is not a fully conscious process and it eventually comes to an end, but the end is not a kind of assent on our part, and instead lies in the causal relevance of that process to our later actions. If we restrict know-that to beliefs, then the belief is merely a residuum of a longer accumulation of know-how, a know-how of how to act in certain cases. We may say that it is a know-how of how to act as though the content of the belief (the proposition) were true, but this is not to say that there is any necessary assent to such content. The belief, in such a light, acts as a description of the agent's behaviour that need not be an element of the agent's conscious experience; it is a second-order description relative to both the object-level of phenomenological experience and the object-level of behaviour (and potential

behaviour), and it is instead an emergent property of these object levels. The belief embedded in a belief–desire model is explanatory of an agent's actions. However, the know-how acts as a precondition for that very belief: the agent pursues its desires as best as it believes, or *knows how*, to.

The belief-in-a-proposition part of know-that is the residuum after all the causally relevant processes of belief formation have finished. Whether or not this process was partly conscious or not does not matter, for it is ultimately rooted in unconscious causal factors that determined how the process was carried out, and the resulting guidance it provides for the agent's actions. The components of that process can be said to have the function of generating that conscious thought or belief; they know how to make you believe. But, the components are also part of the system that does the believing; thus, it can be said that the system itself must know how to believe (in the linguistic sense) before it does believe. This mimics Wittgenstein's use of knowing how in the earlier "it only makes sense to ask what something is called if he already knows how to make use of the name". The belief is in this sense just the name of the thought that led to it and the thought that will lead to further action, but the thought referenced here is an unconscious one, the thought that led to the realization of the conscious thought.[5]

If these conclusions are correct and know-that is a form of know-how, it is still important to spell out how exactly it reduces to know-how. If know-how is teleological in nature and is the building block by which know-that is constructed, then know-that must also be teleological in nature. The reduction, I propose, is a functional decomposition. Knowing that $2 + 2 = 4$ is a matter of knowing how to work with certain mathematical rules, a matter of knowing how to add, how to equate, how to manipulate symbols, and how to compare such symbols.[6] In another light, knowing that there is a tree in front of you is a matter of knowing how to act in relation to the perceived tree. The philosophy of biology may shed further light on the know-how of perceptions through the notion of affordances: organisms perceive affordances and bias their causal repertoire in order to pursue their goals (Fulda 2017, 84). An affordance is an

[5] For a lack of a better vocabulary, I use the term thought for both the conscious and unconscious thoughts. Causally, I don't believe there to be a difference between conscious and unconscious thoughts; it is just the case that only some thoughts clearly present themselves to the consciousness while others remain in the shadows and nevertheless inform our actions through gut feelings or quick intuitions.

[6] Debates concerning the sociology of math and what it means to know a mathematical truth were prominent at the end of the twentieth century, and were often framed through a Wittgensteinian lens. Knowing a mathematical truth, a mathematical operation or other procedure, was equated with knowing how certain rules generalize. For more discussion, see Lynch (1992), Bloor (1992), and Kusch (2004).

ecological property or disposition which has a practical relevance to an agent's possible actions. For instance, a tree can be approachable (in a particular way, with certain rotations, steps, and motions), climbable, choppable, burnable, etc., and knowing that a tree is in front of you just boils down to knowing how to act according to the affordances the tree in front of you has. Crucially, the affordances are framed in relation to the actions the agent can take, and thus related to the causal repertoire of the agent, that is, what it knows how to do. In general, any case of know-that may be decomposed into its relevant know-how through its action-guiding content, be it framed through rules, affordances, or more explicitly know-how.

Let us close with some Wittgenstein: "In one sense knowing is to have learned and not forgotten. In this way, it hangs together with memory" (Wittgenstein 1980, 300). The belief, as the object of assention, is just the conscious representation of the thought that one actually has (or had). It comes with the additional know-how of knowing how to communicate it to others, knowing how to express it to others or to oneself, and knowing how to remember it. But, it is still just a form of that more general knowledge that was learned *and* forgotten—the know-how that guides actions and the pursuit of the agent's goals, irrespective of whether or not it is conscious or assented to.

Conclusion

There is an underlying teleology in existing accounts of know-how. However, all such accounts require references to phenemonological or intellectual entities in their formulation, which pigeonholes know-how as a purely intellectual topic. However, it is possible to form a natural and teleological know-how that is similar in spirit to how philosophers of biology now view natural agency. What is required is a goal-directed system that robustly pursues a goal in contextually relevant environmental conditions. No reference to a cognitive state or representation of the world is required. The proposed teleological know-how is then applied to many old and new examples, showing its flexibility and how it diagnoses previous intuitions. Such a conception of know-how is shown to be applicable more broadly, to cases in social epistemology as well as non-human agents. A teleological account of know-how seems to be a more informative framework with which to approach epistemology.

It is further argued that know-that reduces to know-how. This lends itself to Wittgenstein's later writings that discuss what it means to have meaning in a language, what it means to draw an inference, and what it means to know how. The unconscious know-how of an agent leads them to form their conscious thoughts and consequently their beliefs; beliefs are constructed out

of know-how and functionally decompose into it. In addition to having the unconscious thought, it may be said that know-that is also to know how to express a belief or communicate it. However, it is still know-how all the way up.

Acknowledgments
I would like to thank Joseph Berkovitz for foundational discussions regarding knowledge-how and Denis Walsh for introducing me to the philosophy of biology and teleology. Additionally, I would like to thank an anonymous reviewer for helpful suggestions.

References

Ågren, J. Arvid. 2021. *The Gene's-Eye View of Evolution*. First edition. Oxford Scholarship Online. Oxford: Oxford University Press. https://doi.org/10.1093/oso/9780198862260.001.0001.

Baedke, Jan, and Alejandro Fábregas-Tejeda. 2023. "The Organism in Evolutionary Explanation: From Early Twentieth Century to the Extended Evolutionary Synthesis." In *Evolutionary Biology: Contemporary and Historical Reflections Upon Core Theory*, edited by Thomas E. Dickins and Benjamin J.A. Dickins, 6:121–50. Evolutionary Biology – New Perspectives on Its Development. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-031-22028-9_8.

Bengson, John, and Marc A. Moffett. 2007. "Know-How and Concept Possession." *Philosophical Studies* 136 (1): 31–57. https://doi.org/10.1007/s11098-007-9146-4.

———. 2012a. *Knowing How: Essays on Knowledge, Mind, and Action*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780195389364.001.0001.

———. 2012b. "Nonpropositional Intellectualism." In *Knowing How: Essays on Knowledge, Mind, and Action*, edited by John Bengson and Marc A. Moffett, 0. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780195389364.003.0007.

Bloor, David. 1992. "8. Left and Right Wittgensteinians." In *Science as Practice and Culture*, edited by Andrew Pickering, 266–82. University of Chicago Press. https://doi.org/10.7208/9780226668208-009.

Carroll, Lewis. 1995. "What the Tortoise Said to Achilles." *Mind* 104 (416): 691–93.

Cath, Yuri. 2019. "Knowing How." *Analysis* 79 (3): 487–503. https://doi.org/10.1093/analys/anz027.

Chalmers, David John. 2004. *The Conscious Mind: In Search of a Fundamental Theory*. New York: Oxford University Press.

Desmond, Hugh, and Philippe Huneman. 2020. "The Ontology of Organismic Agency: A Kantian Approach." In *Natural Born Monads*, edited by Andrea Altobrando and Pierfrancesco Biasetti, 33–64. De Gruyter. https://doi.org/10.1515/9783110604665-003.

Ditter, Andreas. 2016. "Why Intellectualism Still Fails." *The Philosophical Quarterly* 66 (264): 500–515. https://doi.org/10.1093/pq/pqv115.

Elzinga, Benjamin. 2021. "Intellectualizing Know How." *Synthese* 198 (2): 1741–60. https://doi.org/10.1007/s11229-019-02160-6.

Fridland, Ellen. 2013. "Problems with Intellectualism." *Philosophical Studies* 165 (3): 879–91. https://doi.org/10.1007/s11098-012-9994-4.

———. 2017. "Automatically Minded." *Synthese* 194 (11): 4337–63. https://doi.org/10.1007/s11229-014-0617-9.

Fulda, Fermin. 2017. "Natural Agency: The Case of Bacterial Cognition." *Journal of the American Philosophical Association* 3 (1): 69–90. https://doi.org/10.1017/apa.2017.5.

Fuller, Gareth. 2022. "A Defence of Functional Kinds: Multiple Realisability and Explanatory Counterfactuals." *International Studies in the Philosophy of Science* 35 (2): 119–33. https://doi.org/10.1080/02698595.2022.2144436.

Gambarotto, Andrea. 2019. "Teleology: A Case Study in iHPS." In *The Past, Present, and Future of Integrated History and Philosophy of Science*. Routledge.

Harrison, Britt. 2013. "The Epistemology of Know-How." Thesis. http://uhra.herts.ac.uk/handle/2299/10433.

Hornsby, Jennifer. 2016. "Intending, Knowing How, Infinitives." *Canadian Journal of Philosophy* 46 (1): 1–17. https://doi.org/10.1080/00455091.2015.1132544.

Jaeger, Johannes. 2024. "The Fourth Perspective: Evolution and Organismal Agency." In *Organization in Biology*, edited by Matteo Mossio, 33:159–86. History, Philosophy and Theory of the Life Sciences. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-031-38968-9_8.

Knorr-Cetina, K. 1999. *Epistemic Cultures: How the Sciences Make Knowledge*. Cambridge, Mass: Harvard University Press.

Kusch, Martin. 2004. "Rule-Scepticism and the Sociology of Scientific Knowledge: The Bloor-Lynch Debate Revisited." *Social Studies of Science* 34 (4): 571–91. https://doi.org/10.1177/0306312704044168.

Laudan, Larry, and Jarrett Leplin. 1991. "Empirical Equivalence and Underdetermination." *The Journal of Philosophy* 88 (9): 449. https://doi.org/10.2307/2026601.

Lynch, Michael. 1992. "7. Extending Wittgenstein: The Pivotal Move from Epistemology to the Sociology of Science." In *Science as Practice and Culture*, edited by Andrew Pickering, 215–65. University of Chicago Press. https://doi.org/10.7208/9780226668208-008.

Nagel, Ernest. 1979. *Teleology Revisited and Other Essays in the Philosophy and History of Science*. The John Dewey Essays in Philosophy 3. New York: Columbia university press.

Noe, A. 2005. "Against Intellectualism." *Analysis* 65 (4): 278–90. https://doi.org/10.1093/analys/65.4.278.

Palermos, S. Orestis, and Deborah P. Tollefsen. 2018. "Group Know-How." In *Socially Extended Epistemology*, edited by J. Adam Carter, Andy Clark, Jesper Kallestrup, S. Orestis Palermos, and Duncan Pritchard, 0. Oxford University Press. https://doi.org/10.1093/oso/9780198801764.003.0007.

Quine, Willard van Orman. 1975. "On Empirically Equivalent Systems of the World." *Erkenntnis* 9 (3): 313–28. https://doi.org/10.1007/bf00178004.

Ramsey, F. P., and G. E. Moore. 1927. "VI.—Symposium: 'Facts and Propositions.'" *Aristotelian Society Supplementary Volume* 7 (1): 153–206. https://doi.org/10.1093/aristoteliansupp/7.1.153.

Ramsey, Frank P. 2016. "Truth and Probability." In *Readings in Formal Epistemology*, edited by Horacio Arló-Costa, Vincent F. Hendricks, and Johan Van Benthem, 21–45. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-319-20451-2_3.

Ryle, Gilbert. 2009. *The Concept of Mind*. Edited by Julia Tanney. London: Routledge.

Schroeder, Mark. 2012. "Showing How to Derive Knowing How." *Philosophy and Phenomenological Research* 85 (3): 746–53. https://doi.org/10.1111/j.1933-1592.2012.00639.x.

Sommerhoff, Gerd. 1950. *Analytical Biology*. Oxford University Press.

Stanley, Jason. 2011. *Know How*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199695362.001.0001.

Stanley, Jason, and Timothy Willlamson. 2001. "Knowing How:" *Journal of Philosophy* 98 (8): 411–44. https://doi.org/10.2307/2678403.

Sultan, Sonia E., Armin P. Moczek, and Denis Walsh. 2022. "Bridging the Explanatory Gaps: What Can We Learn from a Biological Agency Perspective?" *BioEssays* 44 (1): 2100185. https://doi.org/10.1002/bies.202100185.

Uller, Tobias. 2023. "Agency, Goal Orientation, and Evolutionary Explanations." In *Evolution "On Purpose,"* edited by Peter A. Corning, Stuart A. Kauffman, Denis Noble, James A. Shapiro, Richard I. Vane-Wright, and Addy Pross, 325–40. The MIT Press. https://doi.org/10.7551/mitpress/14642.003.0020.

Van Fraassen, Bastiaan Cornelis. 1980. *The Scientific Image*. Clarendon Library of Logic and Philosophy. Oxford: Clarendon press.

Walsh, Denis M., and Gregory Rupik. 2023. "The Agential Perspective: Countermapping the Modern Synthesis." *Evolution & Development* 25 (6): 335–52. https://doi.org/10.1111/ede.12448.

Wittgenstein, Ludwig. 1968. *Philosophical Investigations*. Oxford: Basil Blackwell.

———. 1979. *On Certainty*. Edited by Gertrude Elizabeth Margaret Anscombe and Georg Henrik von Wright. Oxford: B. Blackwell.

———. 1980. *Remarks on the Philosophy of Psychology*. Edited by G. E. M. Anscombe, G. H. von Wright, and Heikki Nyman. Chicago : Oxford: University of Chicago Press ; Basil Blackwell.

Kye Palider
University of Toronto
Institute for the History and Philosophy of Science and Technology
<kye.palider@mail.utoronto.ca>