European Journal of Philosophy

DOI: 10.1111/ejop.12000

Knowing-how: Problems and Considerations

Ellen Fridland

Abstract: In recent years, a debate concerning the nature of knowing-how has emerged between intellectualists who claim that knowledge-how is reducible to knowledge-that and anti-intellectualists who claim that knowledge-how comprises a unique and irreducible knowledge category. The arguments between these two camps have clustered largely around two issues: (1) intellectualists object to Gilbert Ryle's assertion that knowing-how is a kind of ability, and (2) anti-intellectualists take issue with Jason Stanley and Timothy Williamson's positive, intellectualist account of knowing-how. Like most anti-intellectualists, in this paper I will raise objections to Stanley and Williamson's account of knowing-how and also defend the claim that ability is necessary for knowing-how attributions. Unlike most discussions of knowing-how, however, I will return to more Rylean considerations in order to illustrate that any intellectualist account of knowing-how, not simply Stanley and Williamson's preferred variety, will fail because it will be unable to account for fundamental differences in the knowledge required to instantiate an ability and the knowledge involved in propositional thought.

1. Introduction

In recent years, a debate concerning the nature of knowing-how has emerged. This debate has roots in Gilbert Ryle's famous chapter, 'Knowing-how and knowing-that', which appeared in *The Concept of Mind* (Ryle 1949), but has been revived more recently by Jason Stanley and Timothy Williamson in their now seminal paper, 'Knowing How' (Stanley and Williamson 2001).[1] The lines of battle in this debate have been drawn between intellectualists of the Stanley and Williamson variety who claim that knowledge-how is reducible to knowledge-that, and the anti-intellectualists or neo-Ryleans who claim that knowledge-how comprises a unique and irreducible knowledge category.

The arguments between these two camps have clustered largely around two issues: (1) intellectualists object to Ryle's assertion that knowing-how is a relation to an ability, and (2) anti-intellectualists take issue with Stanley and Williamson's positive account of knowing-how.[2] In this paper, I seek to provide a decisive argument against the idea that knowing-how can be cashed out propositionally. As such, I am squarely in the anti-intellectualist camp. Following tradition, I will review the arguments and strategies that are currently under debate. Like most anti-intellectualists, I will raise objections to Stanley and Williamson's account of knowing-how and also defend the notion that ability is necessary for knowing-how attributions. Unlike most discussions of knowledge-how, however, I will

European Journal of Philosophy ••:•• *ISSN 0966-8373 pp.* ••–••

return to more Rylean considerations of the nature of knowing-how. In the concluding sections of this paper, it will be my goal to provide a robust defence of why knowing-how *cannot* be an instance of knowing-that. I will return to the original question concerning the nature of knowing-how in order to illustrate that *any* intellectualist account of knowing-how, not simply Stanley and Williamson's preferred version, will fail because it will be unable to account for fundamental differences in the knowledge required to instantiate ability and the knowledge involved in thinking propositionally.

I proceed as follows: Firstly, I rehearse and respond to Stanley and Williamson's criticism of Ryle (1949) and their alternate proposal. Next, I elucidate the disparity between knowing-how and knowing-that by presenting an instructive example. From this example, three principles follow: (1) propositional knowledge is not sufficient for know-how, (2) differences in knowledge-how are not proportional to differences in propositional knowledge, and (3) knowing-how to *a* requires the ability to *a*, but it does not require the opportunity to instantiate one's ability to *a*. In the final section of this paper, I explicitly argue that knowledge-how is not an instance of practical reasoning, it is not conceptual, not propositional, and its connection to practice must be accounted for by any satisfactory theory of knowledge.

2. Ryle: Criticism and Response

For years, the distinction between knowing-that and knowing-how was accepted and appealed to, but neither questioned nor developed.[3] Recently, however, a debate surrounding Gilbert Ryle's classic distinction has erupted. Jason Stanley and Timothy Williamson's 2001 article 'Knowing How' has been uniquely responsible for spearheading this resurgence. In that article, they attempt both to undermine Ryle's regress argument and provide an alternative theory of knowing-how.

Stanley and Williamson claim that Ryle's regress does not succeed because in order for it to begin, each knowing-how must be an instance of an intentional action. But often, they claim, the contemplation of a proposition is not intentional at all. And when the contemplation of a proposition is not an intentional event, it is not the type of thing that can fuel the regress. Stanley and Williamson's (2001: 414) argument goes like this: Firstly, they characterize Ryle's regress as follows,

(1) If one Fs, one employs knowledge how to F.
(2) If one employs knowledge that p, one contemplates the proposition that p.

> If knowledge-how is a species of knowledge-that, the content of knowledge how to F is, for some φ, the proposition that φ(F). So the assumption for *reductio* is:

RA: knowledge how to F, is knowledge that φ(F).

> Furthermore, let 'C(p)' denote the act of contemplating the proposition that p. Suppose that Hannah F's. By premise (1) Hanna employs knowledge how to F. By RA, Hanna employs knowledge that φ(F). So, by premise (2), Hannah C(φ(F))s. Since C(φ(F)) is an act, we can reapply premise 1, to obtain the conclusion that Hannah knows how to C(φ(F)). By RA, it then follows that Hannah employs that knowledge that φ(C(φ(F))). By premise (2) it follows that Hanna C(φ(C(φ(F))))s and so on.

Next, Stanley and Williamson argue that this regress doesn't get off the ground because, as they rightly point out, not all doings are instances of knowing-how to do. Their example of Hannah digesting food, but not knowing-how to digest food, is a good example of a non-intentional act.[4] They stipulate that not any activity will start a regress, but only *intentional* actions. From here, Stanley and Williamson move on to consider whether the contemplation of a proposition is always an intentional action. In order to do this, they appeal to Carl Ginet (1975), who argues that not all instances of knowing-that are accompanied by the contemplation of a propositional thought. Ginet (1975: 7) states:

> I exercise (or manifest) my knowledge *that* one can get the door open by turning the knob and pushing it (as well as my knowledge *that* there is a door there) by performing that operation quite automatically as I leave the room; and I may do this, of course, without formulating (in my mind or out loud) that proposition or any other relevant proposition.

Ginet's claim, if true, would arguably render premise (2) false. But, beyond insisting that premise (2) is false, Stanley and Williamson choose to argue that contemplating a proposition, in some cases, is not an intentional action at all—that is, contemplating the proposition that *p* is more like Hannah's digesting food than Sally's raising an objection. Thus, they render such actions improper substitutions for F in premise (1). The force of the regress is deflated.

A. Consciousness and Intentionality

There is a lot that is funny about this argument, but I will limit myself to the most immediate problems. First, Ginet's observations about not needing to contemplate a proposition in order to manifest propositional knowledge cannot lead us to conclude anything about the metaphysical nature of contemplating propositions beyond the fact that, phenomenologically, it does not always *seem* to us that we are contemplating a proposition when we are manifesting propositional knowledge. As Alva Noë (2005) points out, the most we get out of Ginet's claim is that there is not always *conscious* contemplation of a proposition whenever knowledge that *p* is manifest. [5] But this phenomenological absence is hardly a decisive position from which to describe the nature of propositional knowledge.

Importantly, if consciousness cannot tell us anything about the presence or absence of an instance of contemplating that *p*, it seems unlikely that it could tell

us anything about whether that contemplation is intentional or not. The fact remains that one is not entitled to the move from 'it does not seem to me that I am contemplating the proposition that *p*' to 'I am not contemplating the proposition that *p*' or to 'the contemplation of the proposition that *p* is not intentional'. And without this move, Ryle's regress stands unharmed.

B. Applying Propositions and Intelligent Predicates

A further problem for Stanley and Williamson's response to Ryle's regress follows from the fact that, on their account, the contemplation, consideration or implementation of the proposition responsible for guiding an action becomes a passive, automatic process like digestion or circulation. But if contemplating, considering or implementing a proposition is like digestion—things that happen to one, but not things that one does—then they cannot be the kinds of things for which an agent can be called intelligent or stupid. After all, we only apply intellectual predicates to agents based on actions for which they are responsible. One cannot inherit ascriptions of intelligence for processes that are not under an agent's control; we are not intelligent because of our subpersonal processes. The fact is that being intelligent requires doing something. We would never claim that someone digests smartly or stupidly; just like we would never say that one circulates one's blood brilliantly. Importantly, if it turns out that the contemplation of this or that proposition that *w* is a way to ϕ grounds the ascription of intellectual predicates to the agent contemplating them, then we must conclude that either contemplating is an intentional action or that we are, as agents, responsible for that contemplation in some other way.[6] Of course, either option would make contemplating importantly unlike digestion.

At this point, of course, the intellectualist would deny that contemplating a proposition justifies the ascription of intellectual predicates to the agent contemplating it. The first option that the intellectualist has is to claim that it is the presence of a proposition and not its contemplation, consideration, implementation or selection that determines the intelligence of an action and thus justifies the ascription of intellectual predicates to the agent. This would mean that the contemplation, consideration and implementation of the proposition would be irrelevant for determining the intelligent or unintelligent nature of the action.[7]

Such an option proves untenable, however, because it turns out that it is *not* simply the presence of a proposition, but the *appropriateness* of that proposition that determines whether we call an action superb or stupid, smart or silly. It turns out that to be intelligent or unintelligent, we have to consider more than if some proposition is guiding an action; we must consider whether it is the right proposition guiding the action. For instance, let's think of two agents swinging a backhand in the same way. In intellectualist terms, we would cash this out as two agents whose backhand swinging is being governed by the same proposition. However, despite the identity of propositions, it is still possible that one agent swings superbly and the other stupidly. This is because one may swing in

the right circumstances, with the right timing, force and so on, and the other may do it all wrong. For the second agent, she may be governed by a proposition, but it might be the wrong proposition to be governed by given her situation. Importantly, this isn't a unique case where appropriateness matters. The necessity of appropriateness for the attribution of intelligence extends to absolutely every instance of skill that we can think of. However, these considerations show that it is not simply the presence of a proposition that makes an action intelligent, but the presence of the appropriate proposition. If this is the case, then we must ask what could be responsible for an action being governed by the appropriate proposition?

It would seem that for the intellectualist, the mechanism responsible for an action being governed by the right proposition would play the exact same functional role as contemplation does for Ryle.[8] Crucially, for the intellectualist, no matter what she calls the mechanism, some mechanism must occupy this functional role; she needs something to select and apply propositions.[9] Moreover, the mechanism that is responsible for selecting the appropriate propositions will be the mechanism that determines if an action is stupid, silly, smart or brilliant. It will be this mechanism, after all, that determines if the appropriate proposition is governing an action. As such, it turns out that the mechanism that determines the presence of this or that proposition for governing action justifies attributions of intelligence to the agent. Unfortunately, as I pointed out above, if any state, event or process justifies attributions of intelligence to an agent, then it cannot be something that the agent bears no control over. The mechanism cannot be automatic or non-intentional like digestion or circulation. It seems that Stanley and Williamson are wrong.

A second option that the intellectualist has is to say that it is not the mechanism responsible for the presence of this or that proposition that determines the intelligence of an action, but rather the mechanism's relationship to an intentional state that provides the grounds for the attribution of intelligence. The intellectualist could claim that the reason why digestion is not subject to ascriptions of intelligence, but the passive selection mechanisms in action are, is because the latter are connected to intentional states while the former are not.[10] Unfortunately, this option also fails.

Why it fails becomes clear if we notice that we ascribe intellectual predicates to an agent not only based on her having an intention to act, but also based on the fact that the execution of her action proceeds in a certain way or fashion. The point is that intelligence ascriptions that are justified on the basis of the presence of an intention alone are not sufficient to account for the intellectual predicates that are ascribed to agents based on the way or manner in which those intentions are instantiated in action. Importantly, in paradigm cases of skill and intelligent action, we attribute intelligence to agents throughout a performance and not simply at its instigation.

To elucidate this point, let us begin by considering an example from visual perception. Suppose that I intentionally look to my left. I am responsible for looking or attending to a certain area, and thus can be said to look smartly or

obliviously. However, the subsequent process of building up a visual representation is not something that I, the agent, can be assigned normative predicates based upon. If we think of what happens after attending—that is, if we think of the process from retinal stimulation, through the optic nerve, to the visual cortex and so on—that process is not something that can characterize me.[11] The process is at the wrong level to be something for which I can be assigned credit: it is subpersonal, and not personal.

We might say I have a keen visual system, but we wouldn't say that I am keen in the way I process visual information.[12] The normative descriptions that we attribute to the automatic mechanisms in visual perception do not get transferred to me since, for such transference, I must be in control of these mechanisms; I must be responsible for them in some way. As such, we see that looking left intentionally does not entail that the instantiation of the visual process inherits its intelligence from my intention to look. But such inheritance would have to occur if the proposed response on behalf of the intellectualist could work.

Notice how this is different in the case of know-how. When we watch the tennis player Roger Federer swing a backhand, we ascribe to him intellectual states in choosing to swing the backhand—that is, in his intentionally swinging—but we also ascribe intelligence to him based on the way in which he swings. We say he swings 'brilliantly', 'appropriately' or 'beautifully'. We don't say that he 'intentionally swings', and then that the mechanisms that govern his swinging are 'appropriate', 'beautiful' and so on. We apply the predicates directly to Federer because we take it that he is not only the instigator of his action, but the controller of his action throughout its execution. But if this is so, if an agent governs or controls her skills beyond her intention for instantiating them, then we won't be able to claim that intelligence attributions are determined by the intention for which the actions is performed. There will be more to the intelligence of an action then that action being intentional. So, again, we have to ask about how agents can be accountable for their skills if it is really just subpersonal mechanisms that are doing all the important intellectual work.

The take-home point is this: an action being governed by one proposition rather than another cannot be automatic or mechanical since we only attribute intellectual predicates to agents based on the activities for which they are responsible. It follows that if the way a skill is manifest happens to be appropriate as a result of some passive, automatic process, we wouldn't say that the agent 'ϕs appropriately', but that the agent 'intends to ϕ', and then some subpersonal mechanisms appropriately responds to that intention. But clearly, when we call the performance of a musician or an athlete 'brilliant', 'stunning' or 'banal', we are talking about the musician or the athlete and not the passive mechanisms that occur inside the agent—that is, that which determines the suitability of normative intellectual ascriptions to agents cannot be a brute, passive process because otherwise we could only ascribe intelligence to the process, but not to the agent.

These considerations taken together—that there must be a mechanism that is responsible for selecting and applying the appropriate proposition governing

any skilled action, and that the way that mechanism functions justifies our ascription of intellectual predicates to an agent—entails that that mechanism cannot be passive or automatic. As anti-intellectualists we have two options for getting the regress going based on these considerations: (1) we can claim that contemplating a proposition is intentional, or (2) we can admit that contemplating a proposition is not intentional, but insist that it is still intelligent. Then we can claim that the proper substitutions for premise (1) 'If one F's, one employs knowledge how to F' are not limited to intentional actions, but open to all intelligent actions.[13]

C. Learning and PMOPs

In addition to failing to stop the regress, Stanley and Williamson also fail to explain *the* central notion of their proposed intellectualist alternative. Stanley and Williamson (2001: 430) offer the following as a proper construal of knowing-how:

> Hannah knows [PRO how to ride a bicycle] is true relative to a context *c* if and only if there is some contextually relevant way *w* such that Hannah stands in the knowledge-that relation to the Russellian proposition that *w* is a way for Hannah to ride a bicycle, and Hannah entertains this proposition under a practical mode of presentation.

Importantly, on this account, it is *not* simply the content of the proposition that *w* is a way for Hannah to ride a bicycle that does the heavy lifting, but rather, the *way* that this proposition is represented. It is, after all, the practical mode of presentation (PMOP) and not the proposition that is responsible for the difference between someone who *knows-that* ϕ and someone who *knows-how* to ϕ. Stanley (2011a: 5) explains this requirement by saying that:

> It is not enough that there is a way *w* that is a way in which John could find coffee in New York City, and John knows that *w* is a way in which he could find coffee in New York City . . . he must think of the way of finding coffee in New York City *in the right way,* that is, *under the right mode of presentation*. (Italics in original)

As such, the practical mode of presentation shoulders the great majority of the burden for Stanley and Williamson's explanation of knowing-how. It is the practical mode of presentation that accounts for the ability to put knowledge into action, when ability is present. However, Stanley and Williamson offer no explanation to help us understand how the practical mode of presentation accomplishes this task, over and above a rather shoddy analogy with first- and third-person modes of presentation.

Such an account, however, is absolutely central to explaining how knowledge of the proposition that *w* is a way for Hannah to ride her bicycle can be applied in action. And the issue of application is at the heart of the problem of knowing-how. The fact is that though Stanley and Williamson deny that

knowing-how is a relation to an ability, their intellectualist story still needs to account for how knowledge is applied in action when that knowledge *is* applied in action—that is, even if knowledge-how is not identical to ability, this does not mean that an account of knowledge-how can get away with not explaining how knowledge is applied in action when ability is manifest. And this is exactly what Stanley and Williamson leave out by not providing a detailed account of the practical mode of presentation.[14]

Further, we should notice that we are presented with a glaring dissimilarity when we compare learning propositional knowledge with learning knowledge-how. When we learn propositional knowledge, we learn the proposition that *p*. However, on Stanley and Williamson's account, learning knowledge-how does not necessarily centre on learning the proposition that *w* is a way to ϕ since we might already know *that* proposition but just not represent it under a practical mode of presentation. As such, acquiring knowledge-how would involve developing the capacity to represent that proposition under a practical mode of presentation. Oddly, on the intellectualist account, we might spend years practicing to acquire a practical mode of presentation, but never learn anything new. This is because for the intellectualist, all knowledge is propositional, but a practical mode of presentation is not, and so, acquiring this mode of presentation could not be an instance of learning.

However, it certainly seems that when one practices, e.g., playing the piano, one learns something new. It is absurd to claim that when one already knows the guiding proposition 'that *w* is a way to Φ under a non-practical mode of presentation', then what one acquires through skill learning is something brute or noncogntive. So, even if we grant that there is a proposition that guides all instances of knowing-how, and that that proposition is represented under a practical mode of presentation, it would still be natural to say that since what we learn through practice is the ability to represent the proposition under the practical mode of presentation, then we learn something over and above propositional knowledge. It would seem that learning to represent a proposition under a practical mode of presentation is learning something nonpropositional. But if learning is necessarily of knowledge, then it follows that nonpropositional knowledge is necessarily involved in knowing-how. Again, it seems impossible to ignore the nonpropositional aspects of our epistemic economy.

One could ask, of course, whether learning is always of knowledge. It is possible, after all, that some processes that we call learning might involve the acquisition of dispositions, responses or adaptations that are less than completely intellectual. However, we should notice that opting for this explanation is not at all alluring for the intellectualist. After all, since these short-of-knowledge states result from learning, they have to be sensitive and responsive to meaningful information and experience in such a way that would disqualify them from being classified as brute or noncognitive. These non-knowledge states would have to be integrated into our cognitive economy in an informative, relevant and flexible fashion. So, if the intellectualist admits that the practical mode of presentation is acquired through learning, but denies that this learning involves

the acquisition of knowledge, then the intellectualist has tacitly opened up a middle-ground for cognitive but nonpropositional states. But this is precisely what the intellectualist would like to avoid. After all, admitting the existence of nonconceptual or nonpropositional states of proto-knowledge entails accepting the possibility of cognitive states that are not intellectualist in nature

We should also note that if we consider the effort involved in acquiring a practical mode of presentation, then we have an additional dissimilarity with the other more familiar modes of presentation. After all, there is no effort required to learn to represent a proposition under the first- or third-person mode of presentation. We can represent propositions under different modes, but we do not learn to represent propositions in these different ways in anything like the way that we learn skills.

To close, I should say that my point is not that propositional knowledge is not often or cannot ever be a factor in the intelligent exercise of abilities. It is fairly uncontroversial that many abilities require at least some sort of basic propositional knowledge for their successful instantiation. Rather, the point is that if a complete account of propositional knowledge is given, it will not constitute an exhaustive account of knowing-how. Applying knowledge appropriately is not the same as possessing it.

3. An Instructive Example

Bela Karoli is the world-famous coach of several gold-medal-winning women gymnasts. Bela Karoli knows the rules governing the skills that he teaches his athletes, but he is unable to perform these skills himself. It is not the case that he once could perform these skills, but is too fat or too old to perform them now. He never knew how to perform these skills and neither a good diet nor a good time machine could change that fact.[15] Bela Karoli knows the rules governing the skillful performance of, e.g., a standing layout on beam, and he also knows how to express that knowledge in such a way that his gymnasts can apply it to their own learning and performance. He knows *about* how to perform a standing layout on beam. What he does not know, however, is how to perform a standing layout on beam.[16]

Importantly, the missing link is not just that Bela doesn't know what it's like to perform this skill. We can imagine that he is put into a gymnastics-simulator and so has the experience of feeling what it's like to perform the skills he teaches.[17] Upon exiting the simulator, Bela has all the propositional knowledge he can ever have about standing layouts on beam, he has felt what it is like to perform the skill and, for good measure, let's make him young and fit, too. Still, none of this will amount to him knowing how to perform the skill. There is something more to knowing-how to do a standing layout on beam, and that something else is what Bela lacks. It is this something that Bela has never learned.[18]

Imagine that we ask Bela the following question: 'Do you know how to do a standing layout on beam?' An acceptable response would be: 'I know how to

teach someone how to do a standing layout on beam, but I do not know how to do one myself.' Compare Bela's answer to that of the gymnast Mary Lou Reton. If we asked her 'Do you know how to do a standing layout on beam?', she should answer: 'I know how to do a standing layout on beam, but I cannot do one now, since I am too old and out of shape.' Bela cannot perform the skill because he has never learned how to perform the skill; he has never acquired the necessary knowledge. Mary Lou, on the other hand, cannot instantiate her ability due to the fact that the conditions of her physical body act as an obstacle to her performing it.[19] For Mary Lou, the inability to perform a standing layout on beam is a matter of opportunity; it is the inability to put her ability into practice now.[20] But for Bela, it is a matter of lacking the knowledge-how. He cannot put the ability into practice not because he lacks the opportunity, but because he doesn't know how to perform a standing layout on beam. It is vital for us to differentiate between the knowledge-how that Bela doesn't have and the ability that Mary Lou lacks. Mary Lou lacks the opportunity to instantiate her knowledge-how, while Bela lacks the knowledge-how, *simpliciter*.

Importantly, the difference between Bela and Mary Lou will not be obliterated if we give Bela any more propositional knowledge or a more fit body. And crucially, it is not essential or even likely that Mary Lou possesses anywhere near the amount of expressible propositional knowledge that Bela does. We can easily imagine an athlete being unable to report on the rules that she is following in performing a skill while her coach can eloquently provide such details. So, if we are not already in the grips of an intellectualist account, it would seem that Mary Lou has less propositional knowledge than Bela but she still knows how to do something that Bela does not know how to do. Mary Lou, if young and fit, would be able to perform a standing layout on beam, but Bela, even if young and fit, would not.

We should glean at least three principles from the above example: (1) propositional knowledge is usually not sufficient for ability, (2) differences in expertise are not always proportional to differences in propositional knowledge, and (3) knowing how to *a* requires the ability to *a*, though it does not require the opportunity to instantiate ones ability to *a*. These considerations are meant to illustrate the amount of disparity present between propositional knowledge and knowing-how.

A. Propositional Knowledge is Hardly Ever Sufficient for Ability

The above example should make it clear that propositional knowledge is not always sufficient for ability or know-how. After all, Bela knows all the facts and rules governing the performance of a standing layout on beam, but he still does not know how to do a standing layout on beam.[21] This is because knowledge of rules and facts is not identical to the practical knowledge of how to put those facts and rules into practice effectively.[22] If it were, then anyone that had the requisite knowledge and opportunity could successfully instantiate the ability.

But this clearly is not so. Importantly, what is missing in this case is not more facts or rules.[23] If Bela wants to learn how to perform the skill, he should not consult an encyclopedia, but rather he should practice the balance beam. He should acquire the knowledge-how to perform the skill and this knowledge involves learning the appropriate attention and control required for performing the skill. Significantly, knowing these things is not identical to knowing facts about these things; it is not identical to knowing that to which one ought to attend, knowing what one ought to control, or the way in which one ought to respond.[24] After all, Bela already knows that.

Again, if we agree that learning is always of knowledge, but that learning an ability involves more than simply learning propositions, then we must agree that some part of what is learned when one learns knowledge-how is not knowledge-that.

It may be argued that there are some instances where the possession of propositional knowledge *is* sufficient for ability. Though I am not convinced, I think it will be useful to examine some proposed examples. In 'Propositional Knowledge and Know-how', John Williams (2008) presents the example of Sam who has been told the way to get to the post office. Sam has been given directions, and simply by virtue of being in possession of these new propositional facts, Williams claims that Sam now has the ability to get to the post office; he claims that these facts are sufficient for know-how. I disagree. It seems obvious that simply having the requisite propositional knowledge (that to get to the post office one must, e.g., take a right, walk two blocks, take a left, and then an immediate right) does not guarantee anyone's ability to get to the post office.[25] After all, if it did then this would mean that anyone that was in possession of this propositional knowledge could get to the post office. But, that is simply false. People with good directions get lost all the time.

What if Sam always confuses his right for his left, or he doesn't know which instruction to follow first and which second, or he doesn't know that turning right requires turning right on a street and not an alley or a driveway? It seems that Sam has to know a lot about the world and how to act in it in order to follow directions. That is not to say that none of Sam's knowledge is propositional, but it is to say that it is far from obvious that *all* of it is. Williams' point is that sometimes just having propositional knowledge is enough to guarantee know-how, but surely getting to the post office requires a whole host of other basic capacities that are not obviously or easily reducible to propositional knowledge.

B. Differences in Expertise are not Proportional to Differences in Propositional Knowledge

In the above example, I posit that Mary Lou has less propositional knowledge about how to do a standing layout on beam than Bela but she still knows how to perform the skill while Bela does not. This suggests that expertise in

performance is not proportional to differences in propositional knowledge. It is not always the case that the more one knows about an ability the better one will be at performing it. After all, being a linguist does not entail being a great speaker. And many people speak grammatically without being able to cite the rules of grammar.

In fact, it is often noted that skilled performers regularly fail to describe the rules that govern their performance accurately. Consistently, experts either fail to state or state falsely the rules that govern their behaviour and the cues that they rely on in order to apply these rules.[26] Many skilled performers cannot report on the rules and facts that they actually use to successfully instantiate their abilities. This makes it unlikely that they instantiate their expertise following anything like the directives of propositional knowledge. After all, even if at the time of performance we do not require the conscious contemplation of a proposition in order to attribute propositional knowledge to an agent, it seems odd that someone in possession of the requisite conceptual repertoire and linguistic competency, when prompted, could not reflect and report on her knowledge—that is, even if we admit that knowing how to ϕ and knowing how to explain how to ϕ are different abilities, if propositional knowledge was in fact governing behaviour, what kind of block could experts have in accessing it? And why doesn't this happen with all sorts of other propositional beliefs?

Think for a moment about your own reasoning skills. Is it possible that even with knowledge of different logics, you still cannot cite the rules governing your comprehension of these sentences and the way in which you are coming up with objections to them? Try it. Compare this to reasoning through a problem using formal logic. Wouldn't it be easy for you to state the rules that you are following in that case? It seems that the difference is that in the second instance there are rules that you know and are actually following, while in the first, even if some rule exists that can be described as the rule that you are following, it is not the case that you know that rule and apply it appropriately in order to reason.[27] After all, the fact that behaviour can be described as instantiating a rule does not entail that the agent knows the rule and performs the action by applying it. I can bounce my leg up and down while listening to music and I'm sure that there's a description of my leg movements that can be given in rule form. Still, it is unlikely that in order for me to keep rhythm with my leg, I need to know or apply anything even remotely resembling that rule. A falling elephant may accelerate at a rate of 9.8 meters per second per second but it certainly does not need to know the rate at which it will accelerate in order to accelerate at that rate.

This is especially obvious in cases of small children and animals. There are all sorts of creatures to whom we cannot reasonably attribute the type of sophisticated propositional states that would be necessary for them to possess, if abilities were really just propositional thoughts.[28] Dogs explore their environments and play games and respond to familiar faces without having many of the concepts that are required for those abilities to follow from the application of propositional knowledge. An infant may sway to the melody of a children's song

but why should we think that this requires a concept of 'sway' or 'melody'? As Williams (2008: 108) states:

> Shep the sheepdog has been successfully trained to round up sheep, many of us would feel comfortable in saying not only that it has acquired the ability to roundup sheep but also that it now know how to round up sheep. . . . But no subject can hold beliefs that embody concepts which that subject fails to have, and Shep certainly does not have the concept of entropy. However it might have the concept of sheep.

These examples present instances where ability exists in the absence of the corresponding propositional knowledge. These are instances where propositional knowledge is not necessary for knowing-how. Importantly, if knowing-how were reducible to knowing-that, then we should not be able to have one without the other.[29]

This is not to say that propositional knowledge is not necessary for any ability. Some abilities will most certainly require propositional knowledge for their successful instantiation. And it seems plausible that the more complex the ability, the more propositional knowledge will be required in order to perform it. But, for example, even if I have to know some things about pianos and keys and timing and notes in order to play the piano, it is not the case that my ability simply follows from my knowledge of those facts.

C. Knowledge-how, Ability and Opportunity

It has often been argued that knowing-how cannot require ability since there are cases where knowing-how attributions are warranted, but a subject is unable to perform the ability in question. Carr (1981: 53) states:

> In the first place, although it might seem reasonable to argue that possession of physical ability must be considered a necessary condition of knowing how to do something, it is clear, surely, that there is nothing in the least paradoxical about describing an elderly and arthritic piano teacher or a temporarily incapacitated gymnast as knowing how to do whatever they cannot currently perform . . . present possession of the ability does not appear to be necessary.

Along with the arthritic piano teacher and the incapacitated gymnast, I'd like to say that the amputee cyclist and me, when I'm in a dark room, still have the respective abilities of playing the piano, doing gymnastics, riding a bike and reading. In these cases, we all know how to *a* without being able to *a*, right now.

But does the inability to perform a skill now indicate that success is not a necessary criterion for skill? I argue that the answer to this question is 'no'. For, despite the fact that counterexamples seem to abound, the ability to successfully perform *a is* a prerequisite for knowing-how to *a*. In order to reconcile this

apparently contradictory position, I think we ought to appeal to Katherine Hawley's (2003: 20) account of knowing-how, where she claims that 'actual success is not a necessary condition for knowledge-how, but success under certain counterfactual circumstances is indeed necessary'. What makes ability attributions true is not that one can perform successfully under *any* circumstances, but that one can perform successfully under circumstances that are presumed to be *normal*. When circumstances are not normal, through the use of counterfactuals, we can assess whether a person would, under normal circumstances, be able to instantiate her knowledge-how.

The above examples help us to see that the reason failure doesn't count against a subject's possessing know-how is because the conditions in which the skill is executed are abnormal. We say that the arthritic piano teacher knows how to play the piano because we consider normal conditions for playing the piano having fingers that can move with dexterity. If the piano teacher satisfied those conditions, then indeed he'd be able to play the piano. In the same way, we see that the amputee cyclist knows how to ride a bike because, if he had two legs, he'd be able to ride a bike.[30] And also, I know how to read, even in the dark, because if it were light, I'd be able to read. After all, no one believes that to know how to read requires the ability to read in all circumstances. If I don't have my reading glasses, or the proper lighting or access to reading material, I cannot read, but I still possess the ability to read. Only if normal conditions were met, only if I had the opportunity to read but was not able to read, would we say that I do not know how to read. In fact, often in abnormal conditions, it is quite natural to say that *S* has the ability to *a* even if *S* is not able to *a*.

For every ability there are opportunity conditions that we assume must be satisfied in order for failure to *a* to count against ability possession or know-how. If these opportunity conditions are not met, then failure to *a* does not preclude having the ability to *a*. Rather, we should understand such cases as lacking the opportunity to instantiate the ability to *a*. As Hawley (2003: 22) states: '[F]ailure to perform in abnormal circumstances for that task doesn't usually count against someone's possessing knowledge-how.' Of course, 'normal' conditions are notoriously difficult to specify and though it would be convenient to avoid this kind of talk, I do not think that it is at all possible. The closest that we can get to an accurate account of know-how will involve the messy and complicated conditions of the world.

Though 'normal conditions' are messy they need not be arbitrary—that is, we can place limits on normal conditions in a way that makes them informative. In order to do so, we must address at least the following problem: it may seem that we can present normal opportunity conditions in such a way that everyone knows how to do everything. So, we can say that I know how to fly a plane because if I went to flight school and practiced flying for 80 plus hours and so on, then I'd be able to fly a plane successfully. However, it would be rather odd to build the learning conditions for a skill into the normal opportunity conditions for possessing it. As such, I think we have a natural limitation on the types of counterfactual circumstances that we ought to bring to bear in determining

normal counterfactual circumstances. Quite clearly, any sort of learning, training, drilling or developing conditions cannot be legitimate posits for assessing opportunity conditions. After all, we want to determine which skills agents really possess, not the skills that they might possess, if they were to learn them.

Returning to the previous example, remember Bela who has oodles of propositional knowledge but does not possess know-how and Mary Lou who has less propositional knowledge than Bela but possesses the knowledge-how that Bela lacks. Recall that neither Bela nor Mary Lou can presently perform a standing layout on beam. We can now explain why we attribute know-how to Mary Lou but not to Bela. Mary Lou has know-how because she *would* be able to successfully perform a standing layout on beam if conditions were normal—that is, if Mary Lou met the opportunity conditions of weight, strength and flexibility, then she'd be able to perform the standing layout on beam. On the other hand, even normal circumstances for performing a standing layout on beam cannot provide Bela with the ability to perform successfully. Thus, Mary Lou knows how, but Bela does not.

The considerations from the previous three sections should encourage us to resist the identification between skill and propositional thought by noticing important differences between these two categories. The above alone is not meant to function as an argument against the reducibility of knowledge-how to knowledge-that. Taken with the reflections below, however, I hope to show that cashing out knowing-how in terms of knowing-that is theoretically impermissible.

4. What Knowing-how is Not

Many anti-intellectualist arguments concerning knowledge-how end by pointing out the flaws in Stanley and Williamson's positive thesis. I'd like to do more: I'd like to elucidate why *any* intellectualist explanation will fail to give a satisfactory account of knowledge-how by highlighting some key differences between propositional thought and ability.

A. Know-how is Not Theoretical or Practical Reasoning

It is almost universally accepted that knowledge-how is not theoretical knowledge or reasoning. This is because knowledge-how applies to the realm of action and practice, and theoretical reasoning does not.[31] Though theoretical reasoning can produce deductively valid conclusions, these conclusions are not instructions for action. And while theoretical knowledge can tell us what is true, it cannot tell us how to translate this truth into successful practice. For example, theoretical reasoning can tell us that (1) all horses are animals, that (2) all animals need food to survive, and so, (3) all horses need food to survive, but it cannot tell us that we ought to feed the horses. According to Aristotle, these are the responsibilities of practical reason.[32]

As opposed to theoretical reason, practical reason is reasoning about action. It is instrumental, means–ends reasoning about how to reach one's goals. As such, it seems much more likely that practical reasoning will be able to account for know-how. After all, if practical reason can tell us what to do given particular desires, and know-how is the ability to do things, then at least we are playing the same game here. Practical reason, for instance, can tell us that: (1) if one wants to run a marathon, then one must train, and since (2) I want to run a marathon it follows that (3) I must train. Without some obstacle standing in my way, practical reason should spur my training.

Importantly, the conclusion of this practical syllogism is not an instruction on how to train, but simply the imperative to do so. But we need more than this to account for know-how; we want to know both *that* I ought to train and *how* I ought to train myself, in the middle of winter, with only three months left before the marathon and with this stiff knee, to boot. One may think that if we just get more specific then we can get the practical syllogism to account for the nuances of successful action in particular circumstances, but the *how* of the practical syllogism is too weak to act as a guide for implementing behaviour. After all, successfully performed ability must be responsive to the subtleties and nuances of the exact situation in which one is performing. However, the universal or major premise of the practical syllogism cannot apply to particular situations. If it did, it would cease to be a universal premise.

The more specific one is with the details of the major premise of a practical syllogism, the less it seems one's reasoning is going from the general to the particular. But the less particular one is with one's reasoning, the less likely it is that that reasoning can be responsible for one's ability to successfully perform a skill. Because actions must be executed in particular settings with an acute sensitivity and responsiveness to the actual nuances of this very situation here and now, the major premise of a practical syllogism would have to incorporate all of the very particular details of this very situation, here and now, in order to be responsible for guiding action. Yet this sort of detail undermines the purpose of having general rules from which to derive practical directives.

For example, let us say that I would like practical reasoning to direct me in playing a note on the piano. We'd have to start with something like 'when one wants to play a note on the piano, then one ought to depress a piano key'. But this is not enough to guide behaviour, for we need to know what to depress the piano key with, and so, we need to also say that 'when one want to depress a piano key, one ought to use one's finger'. But then we have to say what part of one's finger one should use to depress the key, and so we have to say that 'when one wants to use one's finger to depress a piano key, one ought to use one's finger tip, at such and such an angle'. But then we should want to know how hard to press, and so, we'd have to say that 'when one wants to depress a piano key with one's finger tip at such and such an angle, then one must press down with the appropriate amount of pressure'. Surely, though, we need to know what the appropriate amount of pressure is, and in order to find this out we must ask

what kind of piano I am playing, and also we should find out when it was tuned last. And from this, we'd have to construct another more detailed premise. But this still isn't enough, for what if I want to play the note loudly? Then we must ask how big the room is, and its shape, and we must find out what is covering the walls. Yet surely all of this cannot be incorporated into the practical syllogism. Or rather, even if it could, what ends up happening is that the practical syllogism doesn't seem to do the work that it is meant to do. This is because to reason by way of the practical syllogism, one must have general rules for how to behave in particular situations. But as it turns out, if one is to use practical reason in order to account for skill, then one does not have general rules, but only particular ones. And problematically, particular rules are not rules at all.

After all, the very concept of a rule contains the idea that it can be applied in multiple situations. For example, the difference between the rule 'when you encounter a '+', take the sum of the two numbers on either side of the sign' and the instruction, 'when you encounter a '+', take the sum of the two numbers on either side of the sign' is that the rule necessarily applies on multiple occasions while the instruction does not. To be applicable on multiple occasions, at least potentially, is exactly what it means to be a rule—that is, this requirement is analytic to the notion of a rule. Further, it is precisely the general character of rules that allows them to play a role in practical reasoning.

However, when considering actions, we must consider them not only in general but in their token instantiations as well. Importantly, we should notice that token manifestations of actions are environmentally embedded in such a way that prevents general rules from governing them—that is, in order to act appropriately in any particular situation one must take into consideration specific circumstances. But taking account of these particular circumstances makes it impossible that general rules are sufficient for directing and guiding behaviour because rules, by their very nature, must abstract from these very particular circumstances.[33]

Aristotle admits this much in Book II of the *Ethics* where he—in no uncertain terms—claims that determining the mean of ethical action is a matter of perception and not reasoning. He states:

> But up to what point and to what extent a man must deviate before he becomes blameworthy is not easy to determine by reasoning, any more than anything else that is perceived by the senses; such things depend on particular facts, and the decision rests with perception. (Aristotle 1988: 1109 a 18)

Aristotle is clear that in practical disciplines like ethics, medicine and navigation, the appropriate action will not be the result of applying a rule. He argues that in circumstances that require particular consideration of situational facts, general reasoning lacks the power to provide a course of action. This is because such reasoning is incapable of taking into account the innumerable variables that

require consideration in order to generate an appropriate response. As such, know-how cannot be the result of practical reasoning.

B. Know-how is Not Conceptual

Furthermore, since skills are concerned with the realm of action and action is always in a particular situation, on a particular material, in a particular environment, we should not expect skills to be constituted by concepts. After all, at the very heart of conceptuality is that idea that concepts can be abstracted away from particular circumstances in order to enter into various relationships while retaining their identity. Concepts are, by definition, general and recombinatorial.[34] But abstracting away from a situation is devastating to the instantiation of an ability since it undermines successful performance: in order for an ability to be successfully performed one must adjust, shift and respond to the very particular circumstances of the environment in which the ability is being instantiated.

Not being sensitive to the particular conditions under which a skill is instantiated sabotages the possibility of that skill's success. For example, if one does a handstand without being responsive to the very material, the very incline and the very uniformity of the surface on which one places one's hands (e.g., a soft mat, or smooth concrete, or a bumpy earthen ground) then one will not be able to perform the micromillimeter, microsecond bodily adjustments required for holding a handstand. The crucial point is this: while concepts can be abstracted away from their environments, skills develop by becoming more and more attuned to their particular circumstances. The movement, as it turns out, is in opposite directions. While concepts are context-*independent*, the elements which constitute skill are very much context-*dependent*.

Importantly, even if there were some context-independent conceptual state that we could posit as that state to which we hold the knowledge relation when we know how to ride a bike or dance ballet, we ought to notice that that general content would never be the content tokened in skill. That content would never be the content responsible for governing a skill. Again, this is because to successfully perform a skill what is needed is sensitivity and responsiveness to context and circumstances, but concepts that meet the generality constraint, by definition, couldn't have this kind of specificity. Crucially, we should also notice how this is different in conceptual *thought*. My concept 'red' is the very concept that shows up when I think 'apples are red' and 'red is a primary colour. Concepts get tokened in thought, but context-independent elements do not get tokened in skill. As such, it is safe to conclude that skills are not conceptual.

C. Know-how is Not Propositional

From the fact that skills are not conceptual, it follows that skills cannot be propositional. This is because propositions are necessarily constituted by con-

cepts. So, though skills may be related to facts, rules, theories and all sorts of propositions, they are not identical to any of them. For all such states must be composed of concepts and concepts are incapable of being particular in the way knowledge-how requires. Since knowledge-how is embedded in an environment, its success is dependent on the capacity to adjust and respond to the very circumstances in which it is located—that is, even if there are general rules governing the way to successfully perform an action (and, of course, there are), such general rules can never provide the necessary directives for acting *en situ*. General rules or facts are often quite beside the point since acting *en situ* requires not the application of universal truths, but the discernment of nuances specific to the context in which one is acting. As such, the identity of knowledge-how is partially constituted by the context in which it is instantiated; knowledge-how does not exist in an abstract realm apart, before or beyond the environment in which it is manifest.

Importantly, even if we think of knowledge-how as a type of procedural knowledge, this in no way entails that one and the same procedure becomes activated when knowledge-how is instantiated. Of course, the instantiation of knowledge-how will depend on cognitive procedures, but it is unlikely that numerically identical procedures are instantiated in different environments in order for an ability to become manifest. However, identity would be required if knowledge-how were propositional; if knowledge-how to do a cartwheel was equivalent to knowing-that *w* was a way to do a cartwheel then that very same *w* would need to be represented when one did a cartwheel now, tomorrow or next week.

The heart of the problem is that when an ability is instantiated at one time in one environment and at another time in another environment, it is likely that the procedures that are followed are to some extent created in response to the particular features of the action space. There is no one procedure that is simply turned on every time I do a cartwheel or hit a tennis ball using my backhand. Whereas the proposition 'Sunday is the first day of the week' does not change its structure on Sunday or Monday, in a dirt field or in a pool, abilities, especially embodied abilities, do.[35] To make one's backhand observationally and causally identical when swung on clay and when swung on grass requires differences in the way the backhand is instantiated. To make one's backhand identical on clay and grass requires compensating for the variant features of the surface environment in such a way that the enacted procedures will have to differ from one another in the appropriate ways. It would seem, then, that the more similar embodied skills appear in divergent circumstances, the more variant the procedures that are responsible for their manifestation would have to be. As such, the qualitative identity of an ability in no way entails the numeric identity of the procedures that are instantiated in order to perform it. This means that propositions simply could not do the work needed for the successful instantiation of knowledge-how.[36] After all, a proposition that could only be considered in one context by one person doesn't seem very propositional at all.

D. A Note on Practice

Before closing, it is important to note that any theory of knowing-how should not make the practice of practicing mysterious. Learning knowledge-how is unique in many ways, but the necessity of practice and repetition long after initial success is the hallmark of acquiring practical expertise. Practice requires a type of drilling, training and repetition different from fact-oriented repetition: often, when a skill is practiced, it is repeated until it becomes automatic; it is drilled until it becomes second nature. There does not seem to be a parallel between this sort of practice and the learning of facts or rules. Though memorizing a fact may require repetition, one need not, like with knowledge-how, repeat the fact after one has learned it. The style of a fact that is memorized does not change with repetition: it does not become more true. The manner in which knowledge-how is manifest, however, often depends squarely on the practice by which it is learned.

Given such discrepancies, we should want to know why this particular brand of learning is required for knowledge-how but not for knowledge-that. After all, it seems initially plausible that the method of knowledge-acquisition has some bearing on the nature of the knowledge that is acquired via that method. This is not to claim that every learning method results in a unique kind of knowledge, but it is to suggest that there are enough dissimilarities between learning propositional knowledge and learning knowledge-how to warrant reflection.

Importantly, it is not the case that the learning methods for knowledge-how and knowledge-that are interchangeable. It is not the case that if one wanted, one could study a book instead of practicing on the balance beam in order to learn how to do a standing layout on beam. The necessity of practice for learning knowledge-how must be explained by a theory of knowing-how. Unfortunately, any account of knowledge-how that denies that the knowledge acquired through practice is distinct from that involved in learning facts and rules will be unprepared, in principle, to provide a satisfactory explanation. Moreover, it seems that a theory of knowledge-how, which is committed to the idea that knowledge-how is a version of knowledge-that, will likely not be able to call the results of learning through practice 'knowledge' at all. This is because if the result of learning through practice is not propositional, and if propositional knowledge is the only kind of knowledge, then practice will not yield knowledge but something else entirely.

5. Conclusion

To summarize, whatever know-how turns out to be, it is not theoretical reasoning, it is not practical reasoning, it is not conceptual, it is not propositional and it is learned through practice. I take it that this paper, by presenting a defence of the distinction between knowing-how and knowing-that, provides legitimacy to philosophers who wish to pursue explanations of knowledge, which are not

simply derivative of accounts of propositional thought. I take it that if we can secure a way that may allow us to trace the development of knowledge from nonconceptual, embodied abilities all the way to inferentially promiscuous propositional states, then we are making headway in providing a naturalized account of cognition. I take it that the arguments in this paper lay a foundation for developing such an account.[37]

Ellen Fridland
Berlin School of Mind and Brain
Humboldt University of Berlin
Germany
ellenfridland@gmail.com
ellen.fridland@philosophie.hu-berlin.de

NOTES

[1] Though this chapter from *The Concept of Mind* is the most famous formulation of Ryle's arguments concerning knowing-how, it is equally important to note that the distinction between knowing-how and knowing-that was first discussed by Ryle in his 1946 paper 'Knowing How and Knowing That' (Ryle 1946).

[2] Of course, framing the problem in this way is a simplification. However, for the purposes of a comprehensive overview, especially when it comes to Stanley and Williamson-style Intellectualism, I see these broad brushstrokes as particularly useful. There are, of course, other varieties of Intellectualism, which I will not address directly in this paper. Though I will not address these other varieties directly, my critique in the final section of this paper should apply to Intellectualism, broadly construed. For more on the varieties of Intellectualism and the complexities of this debate, see Bengson and Moffet (2012).

[3] For examples of this from philosophy of mind and language, see Stanley and Williamson (2001).

[4] Alva Noë (2005: 279) argues that to say that digesting food is something Hannah does is suspect in the first place. It is not an act at all, but a happening. Hannah does not *do* anything, well, fast or otherwise.

[5] 'Ryle can accommodate Ginet's observation by countenancing the possibility that not every act of contemplating a proposition is performed consciously. To say that it is or could be performed unconsciously is not to say that it is not the sort of thing that could be performed intentionally' (Noë 2005: 282).

[6] We could follow Davidson (2001: 47) in claiming that intentional actions need not be individuated according to the intentions for them, or alternatively, we could claim that actions and the activities that compose them are intentional in different ways where the former are full-blooded intentional actions while the latter are only under the control or guidance of the agent. Either way, control or intentionality will be at play here.

[7] This is, of course, the option that Stanley (2011b: 26) chooses. He states: 'In the vocabulary of Fodor (1983), triggering representations is something done by an *input system* rather than a *central system*, by a module rather than a central processor. Since triggering representations is something we do automatically, replacing "contemplating a proposition" by "triggering a representation" . . . results in a manifest implausibility.'

[8] For a similar point, see Hetherington (2006: 73).

[9] See Stanley (2011b: 39), citing Fodor (1983): '[K]nowledge doesn't eventuate in behavior in virtue of its propositional content alone. It seems obvious that you need mechanisms to put what you know into action; mechanisms that function to bring the organization of behavior into conformity with the propositional structures that are cognized.'

[10] Thanks to Tobias Rosefeldt for raising this potential objection.

[11] This is the case even if one does not subscribe to a passive notion of perception. Perceptual processing may be active, but it is not, like intentional action, under the control of the agent.

[12] See Noë (2005: 279) for similar points about digestion.

[13] We should note that it was not primarily intentional actions, but intelligent ones that Ryle was concerned with in the first place.

[14] See Williams (2008), Koethe (2002), Noë (2005) and Rosefeldt (2004) for similar points.

[15] I mention this because usually examples of coaches possessing abilities are given in terms of them once having had the ability, but no longer being able to successfully perform it. But the Bela example is different since he never had the ability in the first place.

[16] In Bengson, Moffett and Wright (2009) survey students and find that most would ascribe knowledge-how to someone like Bela. As the above description suggests, however, there is a difference between knowing-how and knowing-*about*-how. This is an ambiguity in our language that, unfortunately, the Bengson, Moffet and Wright survey does not track. See also Rosefeldt (2004) and Snowdon (2003) for assertions of ambiguity in knowing-how uses.

[17] I have in mind here a simulator where Bela is passive and he gets to experience performing the skill without any control or feedback affecting the experience. We could, of course, imagine a simulator where Bela is controlling his movements and, as a result, he may learn how to perform the skill. That is, after all, the point of many simulators (flight and otherwise), but that is more than just getting to feel 'what it's like'.

[18] Stanley and Williamson (2001), of course, would say that Bela lacks a practical mode of presentation.

[19] This possibility of such a case should be uncontroversial since intellectualists agree that there is a difference between knowing how to ϕ and knowing how to explain how to ϕ. See Stanley (2011a: 9).

[20] We can imagine that Mary Lou has a broken ankle if that will help make this point more lucid. This may also help us to abstract from issues of neural plasticity, which seem to implicate the loss of skill. See Noë (2005: 283) for a discussion of this point.

[21] I take it that this example is not question begging since we are all familiar with cases of people that have propositional knowledge but not ability. Importantly, I am not positing some phenomenon that ought to occur *if* propositional knowledge is not sufficient for ability; rather, I am referring to a phenomenon that *does* occur and that must be explained by a theory of knowledge-how.

[22] Of course, Stanley and Williamson (2001) would agree with this. They would say that the knowledge has to be represented in a particular way. My claim is that such a way of representing would itself necessarily be a kind of (nonpropositional) knowledge.

[23] As an example of an abundance of rules without a corresponding ability, Ryle (1949: 41) says the following: '[A] boy is not said to know how to play [chess], if all he can do

is to recite the rules accurately . . . a foreign scholar might not know how to speak grammatical English as well as an English child, for all that he had mastered the theory of English grammar.'

[24] The conflation between knowing how to, e.g., appropriately attend to *q* and knowing the fact that one ought to appropriately attend to *q*, parallels in many ways the conflation between being conscious of *e* and knowing that one is conscious of *e*. The first in no way entails the second, but philosophers are often deluded into thinking that the two claims are identical. See Rosenthal (1991; 1994; 2004) for an explanation of this conflation.

[25] I am ignoring physical obstacles like not being able to walk and not having a wheel chair, or a street being closed because of a parade as the reasons one couldn't get to the post office. I am assuming that one could, if she knew how to follow directions, get to the post-office. What I am questioning, however, is whether physical opportunity plus propositional knowledge is sufficient for knowing-how to get to the post office.

[26] Various examples of this abound in the literature. See Wallis (2008: 130); Berry and Broadbent (1984); Brehmer, Hagafors and Johansson (1980); Reber and Lewis (1977).

[27] One may claim that the difference consists in the fact that in the first instance the rule is consciously contemplated while in the second it is not. This is hardly necessary, though. We can easily imagine a practiced logician not doing anything like consciously contemplating a rule that she is applying, but still being able to cite it as the rule that she is applying, when prompted.

[28] Devitt (forthcoming) forwards a related argument about the implausibility of subjects knowing *the* way to *a*. That is, Devitt claims that possession of a singular concept that *w* is the way to *a* makes it unlikely that subjects are in possession of it.

[29] It is possible, of course, that animals and infants are not in possession of knowledge-how. This is tolerable, but it is not the position that Stanley and Williamson seem to endorse. See Stanley and Williamson (2001: 438).

[30] '[A]fter losing his leg, Seth does not know how to ride a bike under the circumstances of having one leg. But neither did he know how to ride a bike under the circumstances of having one leg *before* he lost his leg' (Williams 2008: Section 7).

[31] See Aristotle (1988, Book VI) for the classic discussion of the distinction between theoretical and practical reasoning.

[32] Also, recall the is/ought distinction in Hume (1982).

[33] When thinking about this problem, I cannot help but be reminded of the movie *Borat* (Charles, 2006) where Sasha-Baron Cohen is being coached on how to make a not-joke. The multiplicity of ways that Cohen misinterprets the instructions just goes to show how nuanced our understanding and sensitivity must be in order to perform even a seemingly simple action successfully. Mr Bean is another good example of how much can go wrong, and so, how much we take for granted that which is necessary for something to go right.

[34] See Evans (1982) for the classical discussion of concepts, recombination and generality.

[35] In fact, empirical evidence seems to support the idea that optimal performance incorporates a degree of flexibility that could not be accounted for concepts or propositions. See Angiers (2010) for a popular discussion of this issue.

[36] I assume that one wouldn't want to say that there is a different proposition represented every time one instantiates the same ability in different circumstances. If one did not require the identity of propositions at different times, it would be curious as to why one would insist on propositions in the first place.

[37] I'd like to thank Michael Devitt, John Bengson, Jason Stanley and Tobias Rosefeldt for many helpful comments on this paper. I'd also like to thank the students in my 'Knowing How and Knowing That' seminar, held at Humboldt University of Berlin, for their engagement, questions and challenges.

REFERENCES

Angiers, N. (2010), 'Seeing the Natural World with a Physicist's Lens', *New York Times*, 1 November, D1.

Aristotle (1988), *The Nicomachean Ethics*. Oxford: Oxford University Press.

Bengson, J. and Moffet, M. (2012), 'Two Conceptions of Mind and Action: Knowing How and the Philosophical Theory of Intelligence', in *Knowing How: Essays on Knowledge, Mind and Action*. New York: Oxford University Press.

Bengson, J., Moffett, M. and Wright, J. (2009), 'The Folk on Knowing How', *Philosophical Studies*, 142: 387–401.

Berry, D. and Broadbent, D. (1984), 'On the Relationship between Task Performance and Associated Verbalizable Knowledge', *Quarterly Journal of Experimental Psychology: A Human Experimental Psychology*, 36: 209–231.

Brehmer, B., Hagafors, R. and Johansson, R. (1980), 'Cognitive Skills in Judgment: Subjects' Ability to Use Information about Weights, Function Forms and Organizing Principles', *Organizational Behavior and Human Performance*, 26: 373–385.

Carr, D. (1981), 'Knowledge in Practice', *American Philosophical Quarterly*, 8: 53–61.

Charles, L. (Director), (2006), *Borat*. New York: 20th Century Fox.

Davidson, D. (2001), *Essays on Actions and Events*. Oxford: Clarendon Press.

Devitt, M. (forthcoming), 'Methodology and the Nature of Knowing How', *Journal of Philosophy*.

Evans, G. (1982), *The Varieties of Reference*. Oxford: Oxford University Press.

Fodor, J. (1983), *The Modularity of Mind: An Essay on Faculty Psychology*. Cambridge, MA: MIT Press.

Ginet, C. (1975), *Knowledge, Perception and Memory*. Amsterdam: D. Reidel.

Hawley, K. (2003), 'Success and Knowledge How', *American Philosophical Quarterly*, 40: 19–31.

Hetherington, S. (2006), 'How to Know (that Knowledge-that is Knowledge-how)', in *Epistemology Futures*. Oxford: Clarendon Press.

Hume, D. (1982), *Treatise of Human Nature*. Amherst, NY: Prometheus Books.

Koethe, J. (2002), 'Comments and Criticism: Stanley and Williamson on Knowing How', *Journal of Philosophy*, 99: 325–328.

Noë, A. (2005), 'Against Intellectualism', *Analysis*, 65: 278–290.

Reber, A. and Lewis, S. (1977), 'Implicit Learning: An Analysis of the Form and Structure of a Body of Tacit Knowledge', *Cognition*, 5: 333–361.

Rosefeldt, T. (2004), 'Is Knowing How Simply a Case of Knowing That?', *Philosophical Investigations*, 27: 370–379.

Rosenthal, D. M. (1991), 'Two Concepts of Consciousness', in *The Nature of Mind*. New York: Oxford University Press.

—— (1994), 'State Consciousness and Transitive Consciousness', *Consciousness and Cognition*, 2: 355–363.

—— (2004), 'Being Conscious of Ourselves', *The Monist*, 87: 161–184.

Ryle, G. (1946), 'Knowing How and Knowing That', *Proceedings of the Aristotelian Society*, 46: 1–16.
—— (1949), *The Concept of Mind*. Chicago, IL: University of Chicago Press.
Snowdon, P. (2003), 'Knowing How and Knowing That: A Distinction Reconsidered', *Proceedings of the Aristotelian Society*, 104: 1–29.
Stanley, J. (2011a). Knowing (How). *Nous*, 45: 207–238.
—— (2011b), *Know How*. Oxford: Oxford University Press.
Stanley, J. and Williamson, T. (2001), 'Knowing How', *Journal of Philosophy*, 98: 411–444.
Wallis, C. (2008), 'Consciousness, Context and Know-how', *Synthese*, 160: 123–153.
Williams, J. N. (2008), 'Propositional Knowledge and Know-how', *Synthese*, 165: 107–125.