

Propositional knowledge and know-how

John N. Williams

Received: 30 September 2006 / Accepted: 9 August 2007
© Springer Science+Business Media B.V. 2007

Abstract This paper is roughly in two parts. The first deals with whether know-how is constituted by propositional knowledge, as discussed primarily by [Gilbert Ryle \(1949\)](#) *The concept of mind*. London: Hutchinson, [Jason Stanley and Timothy Williamson \(2001\)](#). *Knowing how*. *Journal of Philosophy*, 98, pp. 411–444 as well as [Stephen Hetherington \(2006\)](#). *How to know that knowledge-that is knowledge-how*. In S. Hetherington (Ed.) *Epistemology futures*. Oxford: Oxford University Press. The conclusion of this first part is that know-how sometimes does and sometimes does not consist in propositional knowledge. The second part defends an analysis of know-how inspired by [Katherine Hawley \(2003\)](#). *Success and knowledge-how*. *American Philosophical Quarterly*, 40, pp. 19–31, insightful proposal that know-how requires counterfactual success. I conclude by showing how this analysis helps to explain why know-how sometimes does and sometimes does not consist of propositional knowledge.

Keywords Knowing how · Propositional knowledge · Counterfactual success · Ability · Reliable methods · Ryle · Hawley · Hetherington · Williamson · Stanley · Lewis

1 Introduction

In daily life we often say that someone knows how to do something or has know-how. For example, we might say that someone knows how to change a car wheel or that she has the know-how needed to change it. So ‘know-how’ is not just a philosophical term

J. N. Williams (✉)
School of Social Sciences, Singapore Management University,
90 Stamford Road, Singapore 178903, Singapore
e-mail: johnwilliams@smu.edu.sg

of art. Attempts to elucidate know-how have been largely overshadowed by analyses of propositional knowledge. Still, an analysis of know-how is worthwhile, not only because it might illuminate other important philosophical debates¹, but also because it is worthwhile understanding a form of knowledge that is important in its own right. After all, it is difficult to see how we could explain our continued survival as a species without mentioning know-how.

In order to obtain a many-sided diet of examples, note that we might say the following: Stella knows how to change a car wheel although she has never tried to change one. After being told, by a reliable informant, how to get to the post office, Sam knows how to get there. Seth the champion cyclist who has just lost his right leg in a car smash, still knows how to ride a bike. Sid knows how to tell the sex of a day-old chick although he has no idea how he does so. By a stroke of luck, Sally has just survived an avalanche, although she does not know how to survive one. Sean does not know how to digest food and Steve doesn't know how to produce white blood cells.

Our intuitions about other attributions of know-how are more labile. On learning that 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 round up sheep but also that it now knows how to round up sheep. Further down the phylogenetic scale, we would all be happy to say that an earthworm has the ability to tunnel through soil, but many of us would judge it inappropriate to say that it knows how to do so. Or to go off the scale to fairly simple inanimate objects, although we might be prepared to say that a thermometer has the ability to register changes in temperature, many of us would not want to say that it knows how to register them. It seems that not all abilities count as pieces of know-how (as I will show in Sect. 2).

This set of intuitions is mirrored by intuitions about attributions of propositional knowledge. Although it is appropriate to attribute propositional knowledge to people, it is less clear that we may attribute it others. Simple inanimate objects (as opposed at least to computers) have no propositional knowledge. It is unlikely that earthworms have it. But higher animals might have rudimentary propositional knowledge. Shep cannot know that entropy is increasing, because (on nearly all accounts of propositional knowledge) if it knows this then it believes it too. But as Searle argues (1992, pp. 155–162), 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, especially since it has been trained to discriminate sheep from non-sheep. So perhaps it may know, for example, that sheep are in the field. It is difficult to know what propositional knowledge we may attribute to Shep because it is difficult to specify, using the linguistic expressions of our thoughts, exactly what concepts are available to it.

¹ It might help us evaluate Lewis's reply to Jackson's attack on physicalism. Jackson (1982, p. 130; 1986, p. 291) supposes that Mary is a brilliant scientist who has spent her life in a black and white room watching black and white television images. Mary comes to know everything there is to know about the physics of colour. When Mary leaves her room and sees red for the first time, she seems to acquire new knowledge. Jackson concludes that Mary has acquired new propositional knowledge about the non-physical world. But Lewis (1983, p. 131; 1990, p. 516) avoids this conclusion by claiming that Mary has come to know how to recognize, remember and imagine experiences of red.

This paper is roughly in two parts. The first part deals with the question of whether know-how is constituted by propositional knowledge, as discussed primarily by Gilbert Ryle (1949), Jason Stanley and Timothy Williamson (2001) and Hetherington (2006). The second part defends an analysis of know-how inspired by Katherine Hawley's (2003) insightful proposal that know-how requires counterfactual success.

In Sect. 2 I argue that know-how isn't actual ability and fault Ryle's analysis of know-how. In Sect. 3 I examine his 'infinite regress' argument that know-how does not always consist of propositional knowledge, as well as Stanley and Williamson's analysis of know-how, which has it that know-how *does* always consist of propositional knowledge. Then I look at their attempt to refute Ryle. I conclude that Hetherington is correct to claim that Stanley and Williamson have not refuted Ryle. In Sect. 4 I examine two of Ryle's lesser-known arguments, that know-how never consists—and in the second argument, does not always consist—of propositional knowledge. The first of these backfires, and in fact makes a reasonable case that know-how *sometimes* consists of propositional knowledge. But his second argument is plausible. I conclude in Sect. 5 by summarizing good objections to Stanley and Williamson's analysis of know-how and adding a couple of my own.

In the second part I defend and complete Hawley's partial analysis of know-how as counterfactual success in completing tasks. In Sect. 6 I argue, very much in her spirit, that we may think of what we know how to do as a task relativised to circumstances individuated by conversational context. In Sect. 7 I defend Hawley's proposal that know-how involves counterfactual success. In Sect. 8 I propose that one's know-how is counterfactual success that one enjoys because of one's possession of reliable methods that one has justification for thinking will result in success. After defending this completed analysis against objections in Sect. 9, I conclude by showing in Sect. 10 how it helps to explain why know-how sometimes does and sometimes does not consist of propositional knowledge.

2 Know-how isn't actual ability

Ryle (1948, p. 33) Lewis (1990, p. 516) and von Wright (1963, p. 48) hold that

$$\forall s \forall \varphi (s \text{ knows how to } \varphi \text{ just in case she has the actual ability to } \varphi).$$

This analysis fails because it is possible that *s* still knows how to perform a task she is newly unable ever again to perform. So knowing how to φ does not entail the actual ability to φ . Consider Seth, a champion cyclist who has just lost his leg in an accident (compare Hawley 2003, pp. 22–23). Evidence that he still knows how to ride a bike is that he has forgotten nothing about riding a bike. Moreover, he could teach others how to ride a bike, and it is plausible that just as you can only teach others propositional knowledge if you have it yourself, so you can only teach others how to do something if you know how to do it yourself.²

² A case in which you do not know how to ride a bike, but only provide a bike, bandages and sympathy, in order to get someone else to know how to ride a bike, is no counterexample because it is not a case of teaching, but only a case of providing opportunities to learn.

Another reason why the analysis fails is that unlike ability, know-how is opaque. Know-how is opaque just in case

$$\diamond \exists s \exists \varphi \exists \chi (s \text{ knows how to } \varphi \text{ yet does not know how to } \chi, \text{ although } \varphi\text{-ing is necessarily equivalent to } \chi\text{-ing}).^3$$

For example, Lois may know how to contact Clark Kent (she knows that she has his telephone number) yet not know how to contact Superman, despite the fact that contacting Clark is necessarily contacting Superman (because Clark is necessarily Superman). Or consider Stan, whose job involves selecting three equal lengths of wood and then gluing them together to make equilateral triangles. Stan, who is not very bright, has the concept of equal length but has no clue what an angle is. He may know how to make equilateral triangles (by following the method above) yet not know how to make equiangular triangles (he has no idea what these are), despite the fact that making equilateral triangles is necessarily making equiangular triangles.

By contrast, abilities are not opaque:

$$\sim \diamond \exists s \exists \varphi \exists \chi (s \text{ knows has the ability to } \varphi \text{ yet does not have the ability to } \chi, \text{ although } \varphi\text{-ing is necessarily equivalent to } \chi\text{-ing}).$$

Since Lois has the ability to contact Clark, she has the ability to contact Superman, although she might not know that she has the latter ability.⁴ Likewise Stan has the ability to make equiangular triangles (by making equilateral triangles) although he does not know that making equilateral triangles is a method of making equiangular triangles.

Since what is true of actual ability is not true of know-how, know-how is not actual ability.

Ryle's variant of the analysis is vulnerable to an extra objection as well. For Ryle, *s*'s actual ability to φ is a complex of dispositions (1949, p. 33). These dispositions are for *s* 'to be bound or liable to be in a particular state, or undergo a particular change, when a particular condition is realized' (1949, p. 43). In other words, such dispositions are expressible as a conjunction of conditionals each of which describes how *s* would definitely or probably behave under specified circumstances. Moreover these

³ As a case of the opacity of know-how, Hawley (2004, p. 26) adopts Davidson's (1963) early example, in which someone knows how to flip the switch but not how to alert a prowler, despite the fact that to flip the switch is to alert the prowler. But opacity must involve substitutions of necessary equivalence. So this example does not establish the opacity of know-how since it only involves substitution of contingent effect for actual cause. Moreover, the opacity of know-how must involve substitutions of the same type of action, but flipping the switch is a different type of action from alerting a prowler.

⁴ Carr (1979, p. 406) claims without argument that from Geach's (1957, p. 15) principle that

If $s\varphi$'s then *s* has the ability to φ

it 'clearly follows' that

If φ -ing is necessarily equivalent to χ -ing, then *s* has the ability to φ just in case she has the ability to χ .

dispositions are ‘multiple-track’ meaning that their actualizations are ‘indefinitely heterogeneous’ or in other words ‘can take a wide and perhaps unlimited variety of shapes’ (1949, p. 43). Since these circumstances may be specified in any number of ways, both the exact number and identity of conjuncts is indeterminate.

This has the dire consequence that it is impossible to know that anyone else has know-how. To know that a conjunction is true is to know that each of its conjuncts is true. But on Ryle’s account, one can never be sure that each of the counterfactual conditionals that conjointly constitute the truth of an ascription of know-how is true, for the simple reason that one can never be in a position to know exactly which are included. This consequence is absurd. Surely we do know that Gary Kasparov knows how to play chess. It would not help Ryle to insist that an attribution of know-how constitutes a *disjunction* of such conditionals. To know that a disjunction is false is to know that each of its disjuncts is false, so now the impossibility becomes that of knowing that anyone else does *not* have know-how. This is equally absurd. Surely we also know that William Hung does not know how to sing.

3 Ryle’s regress argument that know-how does not always consist of propositional knowledge

Ryle wants to show at least that know-how does not always consist of propositional knowledge. He has a number of arguments designed to show this. The one he considered crucial is his ‘regress argument’:

The consideration of propositions is itself an operation the execution of which can be more or less intelligent, less or more stupid. But if, for any operation to be intelligently executed, a prior theoretical operation had first to be performed and performed intelligently, it would be a logical impossibility for anyone ever to break into the circle. (1949, p. 31)

Stanley and Williamson on the other hand, want to show that know-how *does* always consist of propositional knowledge. They start from the linguistic claim that attributions of know-how ‘contain what are called embedded questions’ (2001, p. 418). Roughly, to say that *s* knows how to φ , is to say that she knows an an-

Footnote 4 continued

Carr does not explain why he thinks it follows, although I agree that both principles are highly plausible. Perhaps Carr has the following derivation in mind:

1. <i>s</i> has the ability to φ	Suppose
2. φ -ing is necessarily equivalent to χ -ing	Suppose (all worlds in which <i>s</i> φ 's are identical to those in which she χ 's)
3. <i>s</i> does not have the ability to χ	Suppose for <i>reductio</i>
4. There are close worlds in which <i>s</i> φ 's.	1, plus the principle that it is possible to do what one has the ability to do
5. <i>s</i> does not have the ability to χ in these worlds	3, plus the principle that inabilities spread into close worlds
6. <i>s</i> does not χ in these worlds	5, Geach’s principle
7. <i>s</i> χ 's in these worlds	2,4.
8. <i>s</i> does and does not χ in these worlds	6,7. Contradiction
9. <i>s</i> has the ability to χ	3-8. <i>Reductio ad absurdum</i>

swer to the question ‘What is a way to φ ?’ Based on this, they claim that know-how is a species of propositional knowledge, as captured by the analysis (2001, p. 432):

$\forall s \forall \varphi : s$ knows how to φ just in case

(1) there is some way w which is a way for her to φ

and

(2) she knows that w is a way for her to φ

and

(3) she entertains the proposition that ‘ w is a way for her to φ ’, under a practical mode of presentation.

They represent Ryle’s argument in terms of a *content* that describes a sufficient condition of φ -ing (2001, pp. 413–414). However it turns out that this representation resists translation back into plain English. We should be suspicious of such arguments. So I’ll take the liberty of talking instead of a way w that describes a sufficient condition of φ -ing. Then Stanley and Williamson can be seen as representing Ryle’s *reductio* as

1. $\forall s \forall \varphi$ (If s φ ’s then she employs knowledge how to φ)

2. $\forall s \forall p$ (If s employs knowledge that p then she contemplates ‘ p ’)

Now suppose for *reductio* that

RA) $\forall s \forall \varphi \exists w$ (s knows how to φ just in case she knows that w is a way to φ).

In plain English, RA) says that know-how *always* consists of propositional knowledge.

Suppose that s φ ’s. Then by (1), she employs knowledge how to φ . So by RA, there is a way w such that she knows that w is a way to φ . By (2), she contemplates the proposition ‘ w is a way to φ ’. By (1) again, she employs knowledge how to contemplate the proposition ‘ w is a way to φ ’. So by RA, there is another way w_1 such that she knows that ‘ w_1 is a way to contemplate the proposition ‘ w is a way to φ ’’. And *ad infinitum*. So s must perform infinitely many increasingly complex acts of contemplation before she φ ’s. This she cannot do, but we know that some people do things. So given (1) and (2), RA is false. In other words

$\exists s \exists \varphi \sim \exists w$ (s knows how to φ just in case she knows that w is a way to φ).

In plain English, this says that know-how *does not always* consist of propositional knowledge.

Stanley and Williamson rightly object. (1) is falsified by someone who digests food, because that person does not know how to digest food. Ryle could side-step this example by restricting (1) to intentional φ -ing. Then (2) must be likewise restricted, or else the argument is invalid. But this restriction falsifies (2) because s may φ without contemplating any proposition, as when she manifests, and employs, her knowledge that the phone is ringing by automatically and unreflectively picking up the receiver.

Hetherington (2006, p. 73) complains that Stanley and Williamson have misrepresented Ryle. He quotes a different passage, but one that still talks of intelligent action:

According to the [intellectualist] legend, whenever an agent does anything intelligently, his act is preceded and steered by another internal act of considering a regulative proposition appropriate to his practical problem ... Next, supposing that still to act reasonably I must first perpend the reason for so acting, how am I led to make a suitable application of the reason to the particular situation which my action is to meet?... [T]he absurd assumption made by the intellectualist legend is this, that a performance of any sort inherits all its title to intelligence from some anterior internal operation of planning what to do. (1949, p. 31)

Hetherington concludes that for Ryle, intelligent acting is restricted, not to φ -ing, nor to knowing how to φ , but to φ -ing while knowing how to φ . Based on this, Hetherington constructs a different argument on behalf of Ryle. This also proceeds in terms of a content that describes a sufficient condition of φ -ing. Transposing this into a way w that describes a sufficient condition of φ -ing, it goes as follows. Suppose for *reductio*:

RA') $\forall s \forall \varphi \exists w$ (If s φ 's while knowing how to φ , then she already knows that w is a way to φ , and she knows how to apply this knowledge in order to φ , and she does apply this knowledge in order to φ).

This says, in as plain English as possible, that one's know-how-in-action always involves putting one's knowing-how to apply propositional knowledge to action, into action. Put this way, it is easy to see that a regress follows:

Suppose that s φ 's while knowing how to φ . Then by RA', she already knows that w is a way to φ and she knows how to apply this knowledge in order to φ , and she does apply this knowledge in order to φ . This is a *new* instance of φ -ing while knowing how to φ . Call this φ_1 -ing. So by RA' again, s already knows that w_1 is a way to φ_1 and she knows how to apply w_1 in order to φ_1 , and she does apply w_1 in order to φ_1 . And so on *ad infinitum*. So s must know an infinite number of increasingly complex propositions and she must apply this knowledge at each cycle before she φ 's while knowing how to φ . This she cannot do, but we know that some people do things while knowing how to do them (for example, competent drivers). So RA' is false. In other words

$\exists s \exists \varphi \forall w$ (s φ 's while knowing how to φ & either she does not already know that w is a way to φ , or she does not know how to apply this knowledge in order to φ , or she doesn't apply this knowledge in order to φ).

In as plain English as possible, this says that one's know-how-in-action does not always involve putting one's knowing-how to apply propositional knowledge to action, into action.

Stanley and Williamson's objection to the original argument is ineffective against Hetherington's modification because RA' talks of applying knowledge that w is a way to φ -ing in order to φ . Such an application need not involve contemplating or considering a proposition, but might be unconscious. Hetherington (2006, p. 74) concludes from this that Stanley and Williamson are wrong in claiming that know-how always consists in propositional knowledge.

An objector might observe that all that Hetherington has shown is that

One's know-how-in-action does not always involve putting one's knowing-how to apply propositional knowledge to action, into action.

But what Stanley and Williamson were trying to show was that

One's know-how always consists of one's propositional knowledge.

The objection concludes that since it is possible for both of these claims to be true, Hetherington's case against Stanley and Williamson fails. Stanley and Williamson would say that both claims are true of Seth, because his knowing how to ride a bike consists of propositional knowledge—although because he has lost his leg, he is not able to apply that knowledge to riding a bike.

Hetherington should reply that if RA is the correct analysis of know-how, then RA' must be true of know-how-in-action. But an analysis of know-how that makes it impossible for know-how to be put into action is the wrong analysis. After all, we know that there are people unlike Seth, who do put their know-how into action (for example, working taxi drivers, as opposed to peace-loving gun-owners).

4 Two of Ryle's lesser-known arguments that know-how is not always propositional knowledge

Ryle has four other arguments that know-how is not always propositional knowledge. I'll now look at two of these.⁵ His first is a bad argument that is nonetheless instructive:

⁵ The first of the pair I won't discuss is:

A boy can be said to have partial knowledge of the counties of England, if he knows some of them and does not know others. But he could not be said to have incomplete knowledge of Sussex being an English county. Either he knows this fact or he does not know it. On the other hand, it is proper and normal to speak of a person knowing in part how to do something, i.e. of his having a particular capacity in a limited degree. (1949, pp. 57–58)

This argument has the form

There are cases in which s partially knows how to φ
 There are no cases in which s partially knows that p

s 's knowing how to φ never consists of s 's propositional knowledge about how to φ

This argument is invalid. An opponent may insist that because Sheila knows all the propositional rules of chess except the *en passant* rule, she only partially knows how to make legal chess moves. Yet Sheila has full propositional knowledge of all the rules she does know.

The second of the pair is that 'It makes sense to ask at what moment someone became apprised of truth, but not to ask at what moment someone acquired a skill' (1949, p. 58). This argument has the form:

There is always a definite time at which s learns that p
 There is never a definite time at which s learns how to φ

s 's knowing how to φ never consists of s 's knowledge of facts about how to φ

But the truth of his second premise is dubious. Consider Stephanie, who is learning to ride a bicycle without falling off. She falls off at the first nine attempts but succeeds in staying on at the tenth and always stays on at many subsequent attempts. It seems principled to say that she first learned how to ride a bicycle at her tenth attempt.

... though it is proper to ask for the grounds or reasons for someone's acceptance of a proposition, this question cannot be asked of someone's skill at cards or prudence in investments. (1949, p. 29)

In other words

If s 's knowing how to φ consisted of her propositional knowledge then it would be proper to ask for grounds for knowing how to φ .
But it is never proper to ask for grounds for knowing how to φ .

s 's knowing how to φ never consists of propositional knowledge⁶

The truth of the second premise is a consequence of the fact that since it makes no sense to speak of 'believing how to φ ', it makes no sense to speak of 'grounds for believing how to φ ' either. So it makes no sense to ask for grounds for knowing how to φ .

But the truth of the first premise is questionable. Although s may know that she is in pain or that she seems to see rain, it may not be proper to ask for her grounds for knowing this, because it may not be proper to ask for her grounds for believing that she is in pain or for believing that that she seems to see rain.⁷ More importantly, an opponent may claim that s 's knowing how to φ consists of her knowing facts about how to φ , in other words, consists of her having grounds for beliefs of a kind special enough to constitute propositional knowledge.

For example, Stanley and Williamson may claim that Sam's knowing how to get to the post office consists of Sam's knowing that a way to get to the post office is to leave the main exit, go straight and then turn left. Sam might not have known this to start with. But if a reliable informant were to tell him that this is the way to get to the post office, it seems that Sam has *learned how* to get to the post office and so now knows how to get to the post office. Described one way, as plain know-how, it is not proper to speak of grounds. But described another way, as knowledge of the fact that one way to get to the post office is to leave the main exit, go straight and then turn left, it *is* proper. So s 's grounds attach to her knowing how to φ , but only in the sense that they attach to the beliefs that are an essential part of her body of knowledge of facts about how to φ , a body that constitutes her know-how.

In fact it seems that Ryle's argument backfires: Stanley and Williamson's analysis fits the case of Sam. The right thing to say seems to be that Sam's know-how consists of his knowledge of facts, and that his knowledge is grounded by reliable testimony. So it is reasonable to think that *sometimes*, know-how consists of propositional knowledge.

Ryle has a second lesser-known argument:

It was because Aristotle found himself and others reasoning now intelligently and now stupidly and it was because Izaak Walton found himself and others angling sometimes effectively and sometimes ineffectively that both were able to give their pupils maxims and prescriptions of their arts. It is therefore possible

⁶ I am grateful to a second anonymous reviewer for suggesting a formulation close to this.

⁷ I owe this point to the same reviewer.

for people intelligently to perform some sorts of operations when they are not yet able to consider any propositions enjoining how they should be performed. (1949, p. 31)

This argument has the form

In some cases in which s knows how to φ , s does not have the ability to consider any proposition about how to φ .

In these cases, s knows how to φ , but does not know any fact about how to φ .

Ryle's choice of examples is unfortunate, because they lend his premise no support. For in the sense in which knowledge requires the ability to consider a proposition, it is surely an exaggeration to say that Aristotle or Walton lacked that ability. Had Walton been presented with the claim that 'Worms are good bait for trout' he could certainly have considered it. He had the *ability* to consider it even if in fact he never did consider it.

Nonetheless there seem to be cases in which the premise is true. Shep knows how to round up sheep. But does it have the ability to consider *propositions* about rounding up sheep? The premise might also be true of Sid, who knows how to tell the sex of a day-old chick yet is at a loss to explain how he does it. Perhaps there is a description of ways to tell the sex of day-old chicks, but nobody has ever succeeded in formulating it. This seems like a reason to think that know-how does *not always* consist of propositional knowledge (I will return to this in Sect. 10).

The case of Shep damages Stanley and Williamson's analysis of know-how. Suppose that the only way for it to round up sheep is for it to intimidate all the outlying sheep into a tight flock, and then intimidate the sheep at the rear into moving forward towards the pen. At a stretch we might make a case for Shep having the concept of sheep, but surely it does not have the concept of intimidation. The descriptions of any of the ways in which Shep might round up sheep *could* embody concepts Shep doesn't have. Given Searle's point that this precludes the relevant beliefs, plus the doxastic nature of propositional knowledge, Stanley and Williamson's analysis fails.

5 More problems with Stanley and Williamson's analysis

As we have seen, both Hetherington and Ryle have good objections against Stanley and Williamson's analysis of know-how. John Koethe (2001) also has a good objection: Suppose that Hannah does not know how to ride a bicycle. Susan points to John who is riding a bicycle and says 'That is a way for you to ride a bicycle'. Suppose that this is true. Hannah may now know that *that* way is a way for her to ride a bicycle. But she may still fail to know how to ride a bicycle if she does not know how to instantiate that way herself. So Stanley and Williamson's analysis must include (perhaps replacing talk of 'practical modes of presentation') the condition:

s knows how to instantiate w herself.

This means that Stanley and Williamson must say that if s knows how to φ then she knows, of some way w , that w is a way for her to φ and she knows how to instantiate w herself. But then they must say that she knows of some way w_1 , that w_1 is a way for her to instantiate w and that she knows how to instantiate w_1 herself. This starts a vicious regress.

We might also challenge Stanley and Williamson's fundamental claim that because attributions of know-how contain 'embedded questions', if s knows how to φ then she knows an answer to the question 'What is a way to φ ?' Those who think that it is appropriate to attribute know-how to animals and infants (for example, Baby Sarah who, by following her nose, knows how to find a teat on which to suckle) might claim that such creatures know no answer to any question. Stanley and Williamson could reply that Sarah and Shep may know the answer to a question without knowing that it *is* an answer to a question. But it would be inappropriate for us to *attribute* to Shep even unconscious knowledge of an answer *to a question* because we know that it is in no position to answer any question.

I have a second, even simpler, objection to add: Seth still knows how to ride a bike, despite the loss of his leg recently. But even if he knows of *a way to ride a bike*, there *is* no way *for him* to ride a bike. In fact, normally he will know that *there is no way* for him to ride a bike. This falsifies the first condition of Stanley and Williamson's analysis.

At this point we may draw the moral that an analysis of know-how must respect the fact that how-how sometimes, but not always, consists of propositional knowledge.

In the last part of this paper, I will complete Katherine Hawley's (2003) analysis of know-how as counterfactual success. After defending the completed analysis against objections, I will conclude by showing how this completed analysis helps to explain why know-how sometimes does and sometimes does not consist of propositional knowledge.

6 Individuating tasks by context

In this section, I will follow Hawley (2003, p. 22) very closely, although I will use different examples.

An equivalent way of saying that s knows how to φ is that she knows how to perform the *task* of φ -ing. Although wordier, this formulation allows us to see that tasks come in loose families; loading a shotgun is a different task from loading a rifle, but both may be described as 'loading a gun'. Merely making legal moves on a chessboard is an easier, and thus different, task from legally employing tactics such as pins, skewers, forks, discovered checks and sacrifices, but both may be described as 'playing chess'. Vagueness about which task in its family is in question occurs in at least two ways, either because the task is underspecified, as in 'loading a gun' or because the success conditions of the task vary with conversational context, as in 'playing chess'.

Conversational context helps identify the task we have in mind. Making merely legal moves might count as playing chess when teaching it to small children, but

would not count as playing chess in a chess club. In a British but not a Singapore context it would be reasonable to infer from the fact that Simon knows how to use a washing machine that he knows how to handle hot water. In a military context, but not a civilian one, it is reasonable to infer from the fact that Slim knows how to make a bed that he knows how to fold ‘hospital corners’.

It might be objected that each family of tasks just amounts to a single task (loading a gun or playing chess) that can be performed in different ways (performed on a rifle or a shotgun) or to different standards (requiring legal moves or tactics as well). But since this difference is merely terminological, I see no reason to abandon talk of tasks as belonging to families; such talk helps individuate tasks more finely.

A second objection is that each family amounts to a single task which must be performed to normal, fixed standards. For example, if the standards of success in playing chess (as opposed to playing winning chess) go beyond the making of legal moves, then the novice who only makes legal moves isn’t playing chess, although outside of the chess club we might talk charitably as if she is.

It is true that conversational context often determines what we count as normal circumstances: Stan’s failure to ride a penny farthing doesn’t usually count against his knowing how to ride a bicycle, especially in a club for cycling enthusiasts. But there do not seem to be normal fixed standards for every family of tasks. Rifles and shotguns both provide circumstances for loading a gun and demand different tasks. What it takes to ride a mountain bike downhill over logs is not what it takes to ride a road bike down a firm sandy beach. Although both may be described as riding a bicycle, they are different tasks. Moreover, the presumption that normal circumstances are at stake may be overridden in two ways. First, actual circumstances may be abnormal. For example, we might know that Steve knows how to tie a bowline, yet ask whether he knows how to tie one wearing mittens (the answer might be important on Everest). Second, abnormal circumstances may be explicitly invoked, as when someone claims for a bet to know how to shell a prawn with one hand.

An agent may know how to perform one of the tasks in a family without knowing how to perform other tasks in that family. In other words, she may know how to perform it under some circumstances but not know how to perform it under others. So following Hawley we may say that the claim to be analysed has the form

s knows how to φ , under circumstances *C*.

This should be distinguished from

s knows, under circumstances *C*, how to φ ⁸

for *s* may know, when blindfolded, how to ride a bike, and could do so if she removed the blindfold. In that case the second form is true. But she might not know how to ride a bike when still wearing the blindfold, in which case the first form is false.

⁸ I thank the second reviewer for pointing this out.

7 Know-how involves counterfactual success

Hawley proposes that although knowing how to φ does not require actual success in φ -ing, it does require counterfactual success in the sense that

If s knows how to φ under circumstances C , then if s were to try to φ , under C , then s would succeed in φ -ing.

The fact that someone like Stella may know how to do something she has never tried to do is now easily explained as the fact that the conditional

If s were to try to φ , under C , then s would succeed in φ -ing

may be true despite the actual falsehood of its antecedent.

The conditional also explains why Seth may still know how to ride a bike after the loss of his leg has rendered him unable to ride one. When we judge that he is unable to ride one, we are thinking of present circumstances in which Seth tries to ride a bike with only one leg. But when we judge that Seth still retains his know-how, we are thinking of circumstances in which Seth tries to ride a bike with both legs (Hawley 2003, p. 23). It may still be true after the loss of his leg that

If Seth were to try to ride a bike, under circumstances of having two legs, then he would succeed.

We might be entitled to think, shortly after Seth has lost his leg, that the closest possible worlds to the actual world in which Seth has both legs and tries to ride a bike are worlds in which he succeeds in riding it. Of course, he is physically unable to ride a bike after the loss of his leg. At that time it is false that

If Seth were to try to ride a bike, under circumstances of having one leg, then he would succeed.

After losing his leg, Seth does not know how to ride a bike under the circumstance of having one leg. But neither did he know how to ride a bike under the circumstance of having one leg *before* he lost his leg. So his cycling know-how has not been changed at all by his sudden disability.

None of this means that losing his leg cannot make Seth lose his knowledge of how to ride a bike. Suppose that a sixth-month period of convalescence required by the amputation of his leg has caused his co-ordination to deteriorate dramatically. If we know that this is so, we might be entitled to think that he no longer knows how to ride a bike. But that is because we are entitled to think that if he still had both legs yet was otherwise just as uncoordinated as he is now, his trying to trying to ride a bike would not result in success.

We may now explain the opacity of know-how. Lois knows that she has Clark's telephone number. So if she were to try to contact Clark then she would succeed. Hence Lois knows how to contact Clark. But although contacting Clark is necessarily contacting Superman, *trying* to contact Clark is not necessarily trying to contact Superman. Suppose that the only way of contacting Superman that Lois can think of is to

wait at Metropolis airport where he was last sighted, in the hope of contacting him. But Superman does not know that Lois is trying to contact him and has no intention of returning to Metropolis airport. So if Lois were to try to contact Superman she would fail. She does not know how to contact Superman.

But now a snag seems to arise for cases in which s lacks the ability to try to φ , either because φ -ing is involuntary, as in digesting food or producing white blood cells, or because s is unable to conceive of φ -ing, just as Stan is unable to conceive of making equiangular triangles. In these cases, is it true or false that s knows how to φ ? Stanley and Williamson (2001, p. 440) agree with Ziff (1984, p. 71) that it is false. Given this verdict one might object that

If s were to try to φ , under C , then s would succeed in φ -ing

is true just in case the worlds closest the actual world in the antecedent is true, are worlds in which the consequent is true. But under the circumstances we have in mind, there are *no* worlds in which s tries to φ . So there are no worlds closest to the actual worlds in which s tries to φ . Thus we cannot tell whether the conditional is true. This is absurd, because we do know that it is false that s knows how to φ .

Hawley is cagier:

The ‘trying’ clause doesn’t get off the ground without intentional behaviour: we neither know nor fail to know how to produce white blood cells (2003, p. 26).

Hawley’s point might be that any context in which we ask sensibly ‘Does s know how to φ ?’ is one in which we make the sensible presupposition that s has the ability to try to φ . We might then make a case that in such contexts, ‘Stan knows how to make equiangular triangles’ and ‘Steve knows how to produce white blood cells’ are not just untrue but unintelligible, because the sensible presupposition is false. Or we could argue that there is no fact of the matter of whether s knows how to φ , precisely because the ‘trying’ clause is not truth-evaluable.

Alternatively, we could make it an extra condition of the analysis of know-how that s has the ability, under the relevant circumstances, to try to φ . Then we could agree with Stanley and Williams as well as Ziff that it is false that she knows how to φ . There is no circularity in this extra condition unless s has the ability to try to φ only if she knows how to try to φ . But *trying* to produce white blood cells does not seem like the kind of thing one can know how to do, any more than producing them. It turns out in the next section that there is a much simpler way of resolving this snag.

A more worrying objection is that s may know how to φ , although she is prone to occasional failures of attempts to φ even in ordinary circumstances: even a good cyclist might fall off once in a blue moon, yet still know how to ride a bike. Let us postpone the response to this objection until the next section as well.

8 Know-how is counterfactual success because of possession of reliable methods

Hawley offers three examples (2003, p. 27) to show that counterfactual success is not sufficient for know-how.

Because Sally is prone to mistaking snow for water, she makes swimming motions when hit by an avalanche—which is a good way of surviving an avalanche. If she were to try to survive an avalanche, she would succeed. But she does not know how to survive an avalanche because it is mere serendipity that mistaking snow for water results in survival.

Shelley has a method of making a cake. This is to take whatever ingredients are closest to hand, mix them and bake the resulting mixture.⁹ By lucky chance, the ingredients closest to hand are those that when mixed and baked, result in a cake. If, in these circumstances, she were to try to bake a cake, she would succeed. But she does not know how to make a cake. If the ingredients closest to hand were red wine, cheese and pickled onions instead of milk, butter, eggs and flour, then she would not succeed in making a cake.

The third example goes:

... consider Susie, who likes to annoy Joe, and believes that she does so by smoking. In fact Joe is annoyed by Susie's tapping on her cigarette box, which she does whenever she smokes. Susie would succeed in annoying Joe if she tried, but it seems that she does not know how to annoy Joe, perhaps because she misconstrues the situation. (2003, p. 27)

Hawley's first example is slightly flawed. Surely what Sally would *try* to do is not to survive an *avalanche* for then she would have snow, not water, in mind. However, we can easily repair this by having Sally recognize the snow as snow but base her swimming motions on her misreading a book that in fact gives bad advice about how to survive a flood—bad, because making swimming motions is a reliable method of *perishing* from a flood. Sally's success is still lucky because it is a lucky coincidence that the unreliable method of surviving a flood is also a reliable method of surviving an avalanche.

Incidentally, these cases provide another reason why abilities are not always pieces of know-how; Sally, Shelley and Susie presently have the ability to survive an avalanche, make a cake and annoy Joe, but at no time do they know how to do these things.

Hawley is neutral on the question of what else is needed to define know-how. I propose developing the definition to:

s knows how to φ , under circumstances C , just in case

- (1) if s were to try to φ , under C , then s would usually succeed in φ -ing because
- (2) s has a reliable method of φ -ing, under C .

The 'usually' in the first clause accommodates occasional failure while still explaining the cases of Stella and Seth, as well as opacity.

⁹ Hawley actually says 'throw together' which makes the method even more accidental.

But doesn't the second clause make the first clause redundant? No, because *s*'s possession, in the actual world, of a reliable method of φ -ing is what *explains* her usual success in attempts to φ in nearby worlds: a method is reliable only if it usually results in success. The nature of the explanation—and of any causality involved—is one I will leave aside. Suffice it to say that the first clause is still needed, because *s* may have a reliable method of φ -ing, yet if she were to try to φ , she might, perhaps to amuse herself, use a different, unreliable method of φ -ing that invariably results in failure. For example, Sybil knows how to make a cake, and she has a reliable method of making one, namely following her mother's recipe. But were she to try to make a cake, she would use Shelley's method, with almost inevitable failure. The first clause of the definition blocks this example.

The second clause introduces a kind of externalism into the definition. It blocks the cases of Shelley and Susie. The method that Shelley has of making a cake, which involves taking the closest ingredients to hand, is not a reliable method of making one. A reliable method of φ -ing is a method that results in φ -ing at least more often than not when applied over time in a wide variety of circumstances. In contrast, although Shelly's method would occasionally result in success in lucky circumstances, in most circumstances it would not.

We might argue that the case of Susie is blocked in a different way. What method of irritating Joe does Susie have? Is it the method of smoking or the method of smoking cigarettes from a box which she then taps? This question raises the counterpart of the generalization problem in discussions of epistemic externalism: how narrowly are we to individuate methods in order to assess their reliability? The question is crucial, because although smoking cigarettes from a box which one then taps is a reliable method of irritating Joe, merely smoking is not a reliable method of irritating him, contrary to Susie's belief. A plausible, although not infallible, guide to how we should describe the method that someone has of φ -ing, is how she would describe it herself. At least, this is the default guide we do often use. Since Susie would describe her method of irritating Joe as smoking, that is how we should describe it as well. But this method is unreliable, because in a wide variety of circumstances in which she has cigarettes but no box, it will fail.

However, those think that Susie's actual method of irritating Joe is to smoke cigarettes from a box which she then taps will presumably be prepared to say that she *does* know how to irritate Joe, although she doesn't know how she irritates him. On this view of it, she fits the analysis so far in the same way as Sid.

The second clause also explains why nobody knows how to perform involuntary actions such as digesting food: digesting food is never a matter of method, for methods are intentional, unlike involuntary bodily processes. And Stan is in no position to use any method of making equiangular triangles. This vindicates Ziff's verdict that it is just plain false that Sean, Steve and Stan know how to φ .

The case of Sally is more recalcitrant. Her method of surviving an avalanche is reliable if described as 'making swimming motions' but unreliable if described as 'following the advice of a misread book about floods'. Given that the method that Sally would describe herself as using—making swimming motions—is the method she has, the method she has is reliable, so the definition of know-how is incomplete. This suggests that externalism isn't enough to capture know-how.

Accordingly, we might extend the definition of know-how a final stage further to:

s knows how to φ , under circumstances C , just in case

- (1) if s were to try to φ , under C , then s would usually succeed in φ -ing because
- (2) s has a reliable method of φ -ing, under C that
- (3) s is entitled to believe will usually result in φ -ing.

Note that ‘entitled to believe’ is not ‘believes’ or ‘believes with entitlement’. s is entitled to believe that p just in case if she *were* to believe that p then she would justifiably believe that p . In other words, she ‘has justification’ or has ‘warrant’. This is silent on the question of whether she does in fact believe that p and so is a pretty minimal form of internalism.

This third clause blocks the case of Sally, because she is not entitled to believe that making swimming motions will result in surviving the avalanche; she obtained her reliable method of surviving an avalanche—namely making swimming motions—from a source of reliable methods that is an unreliable source—unreliable not just because it gives bad advice but also because it gives advice about floods, not avalanches.

9 More objections—and more replies

One objection that should be dealt with is that the second condition is too strong for Shep. It knows how to round up sheep, but must have a method of doing so?

Obviously it is unable to produce any description of methods of rounding up sheep. Moreover it is unclear whether an *intention to use* a method is too conceptually sophisticated for sheepdogs. But such intentions are sufficient, not necessary, for it to have a method. Moreover someone’s disposition to describe herself as having a method is not the only form of evidence that she has one. After all, sheepdogs are always trained in ways to round up sheep. So the method that Shep has to round up sheep, is the method that it has been trained to use, despite the fact that it may be unaware of having any method. Moreover there is a sense in which it even has the ability to understand the method it has, for otherwise it could not have been trained in that method. Some of us might even be prepared to say that computers have methods of calculation, even if we balk at the claim that they understand their methods.

Applying the third condition to Shep is more problematic. Can it be entitled to believe that its method of rounding up sheep will usually result in success? Beliefs about the usual success of a method seem far too sophisticated for Shep, even if we make a case for it having an unconscious disposition to form unconscious beliefs. True, the third clause does not oblige us to say that Shep *does* believe that its method of rounding up sheep will usually result in success. It only requires us to say that if it *did* form that belief, then the belief would be justified. But we should admit that this does seem a bit far-fetched as a way of understanding its know-how, because the worlds in which it forms the belief are so far-out. A better response is to simply drop

the third clause in cases of how-how possessed by creatures incapable of beliefs about methods. This marks a principled boundary of most human know-how.

Another objection is that although s knows how to φ , the nearest world in which she tries to φ might easily be one in which circumstances conspire to make her fail to φ . For example, although s knows how to ride a bike, there might be an assassin who never misses waiting to kill her as she tries to mount her bike. I reply that this is hardly the nearest world to the attempt at the task we have in mind; s may know how to ride a bike without knowing how to ride one while being assassinated.

10 Concluding remarks: know-how and propositional knowledge

I have argued that know-how sometimes does, and sometimes does not, consist of propositional knowledge. The Hawley-inspired analysis of know-how helps to explain this.

The analysis easily explains why Sam's know-how consists of propositional knowledge. Sam's knowing how to get to the post office consists of his knowing a set of instructions expressed propositionally. Sam's learning these propositions results in his learning instructions, and once he has learned these, he has counterfactual success because he has the reliable method of following these instructions—one that he is entitled to believe will usually get him to the post office.

At the other boundary we considered two cases in which know-how is not constituted by propositional knowledge, those of Shep and Sid. The analysis of know-how coheres with these. Shep knows how to round up sheep. In other words, if it were to try to round up sheep then it would usually succeed in rounding them up because it has a reliable method of rounding them up, namely intimidation in stages. What explains the fact that this know-how is not constituted by Shep's propositional knowledge is that Shep cannot hold beliefs about intimidation, and this is why in Shep's case the third clause gets dropped.

Sid knows how to tell the sex of a day-old chick. In other words, if he were to try to tell the sex of a day-old chick then he would usually succeed because he has a reliable method of telling its sex that he is entitled to believe will usually tell him its sex. But Sid doesn't know how he tells its sex. He looks at the rear end of the chick and reliably 'sees' what sex it is but has no idea how he makes that decision. It is plausible that there is a reliable method that Sid uses, based on his unconscious recognition of many complex patterns of genital features, but that method is inaccessible to introspection (Horsey 2002). Sid's record of success gives him an inductive entitlement for thinking that whatever he's doing usually works. Yet he is unable to formulate a description of his method. But surely he would have to be able to formulate such a description if his how-how were constituted by propositional knowledge, because then the beliefs that make up that knowledge would form the conceptual basis of the description.

The analysis also explains why our resistance to attributing know-how to creatures lower on the phylogenetic scale goes hand in hand with our resistance to attribute propositional knowledge to them; the point at which we hesitate to speak of a creature

such as an earthworm as trying to do something coincides roughly with the point at which we hesitate to speak of it as having beliefs.¹⁰

References

- Davidson, Donald. (1963). Actions, reasons and causes. *Journal of Philosophy*, 60, 685–700.
- Carr, David. (1979). The logic of knowing how and ability. *Mind*, 88, 394–409.
- Geach, Peter. (1957). *Mental acts*. London: R.K.P.
- Jackson, Frank. (1982). Epiphenomenal qualia. *Philosophical Quarterly*, 32, 127–136.
- Jackson, Frank. (1986). What Mary didn't know. *Journal of Philosophy*, 83, 291–295.
- Hawley, Katherine. (2003). Success and knowledge-how. *American Philosophical Quarterly*, 40, 19–31.
- Hetherington, Stephen. (2006). How to know that knowledge-that is knowledge-how. In S. Hetherington (Ed.) *Epistemology futures*, Oxford: Oxford University Press.
- Horse, Richard. (2002). The art of chicken sexing. *UCL Working Papers in Linguistics*, 14, 107–117.
- Koethe, John. (2001). Stanley and Williamson on knowing how. *Journal of Philosophy*, 99, 325–328.
- Lewis, David. (1983). Postscript to mad pain and martian pain. *Philosophical Papers Volume 1*, New York: Oxford University Press.
- Lewis, David. (1990). What experience teaches. In W. Lycan (Ed.), *Mind and cognition: A reader*, Oxford: Blackwell.
- Ryle, Gilbert. (1949). *The concept of mind*. London: Hutchinson.
- Searle, John. (1992). *The rediscovery of the mind*. Cambridge, Massachusetts: MIT Press.
- Stanley, Jason, & Williamson, Timothy. (2001). Knowing how. *Journal of Philosophy*, 98, 411–444.
- von Wright, & Georg Henrik. (1963). *Norm and action. A logical enquiry*. London: Routledge & Kegan Paul.
- Ziff, Paul. (1984). *Epistemic analysis*. Boston: Reidel.

¹⁰ This paper was written as part of a project funded by a research grant at Singapore Management University (grant number: C208/MSS5E022). I am grateful for correspondence with Katherine Hawley and Stephen Hetherington. Discussion with Tan Yoo Guan, Brian Mooney and Mark Nowacki helped a lot. Two anonymous reviewers gave me useful comments on earlier drafts. The sharp objections of the second of these obliged me to rethink the paper in many ways.