

Success and Knowledge-How

Katherine Hawley, *American Philosophical Quarterly*, 40.1 (2003), 19-31.

1. Introduction¹

Modern epistemologists don't often discuss knowledge-how - propositional knowledge has attracted the lion's share of attention.² Yet the notion of knowledge-how looks useful elsewhere in philosophy - philosophers of science discuss tacit knowledge and skills, philosophers of mind disagree about whether knowing what an experience is like is a matter of knowing how to imagine or recognise it, and philosophers of language and of value consider whether knowledge of meaning or morality is knowledge-how (to use words, to follow rules, to behave well).³ Without a fuller understanding of knowledge-how, it is difficult to assess these proposals. Moreover, knowledge-how is interesting *qua* species of knowledge - when set alongside theories of propositional knowledge, enquiry into other forms of knowledge promises to shed light on the nature of knowledge quite generally. What is it that we value about our interactions with the world, whether these interactions are through belief or through action?

Following Gilbert Ryle (1949),³ those epistemologists who have discussed knowledge-how have often been concerned primarily with the relationship between knowledge-how and propositional knowledge. Most recently, for example, Jason Stanley and Timothy Williamson (2001) have argued that knowledge-how is a certain type of propositional knowledge about ways to act. This paper describes a different approach, one which tackles knowledge-how more directly. Just as propositional knowledge can be understood in terms of true belief plus warrant, knowledge-how can be understood in terms of successful action

plus warrant.⁴ More exactly, success under certain counterfactual circumstances is a necessary condition for knowledge-how - the main part of the paper is devoted to explicating this claim about success.⁵ Two alternative views will be considered and rejected - the claim that success is not necessary for knowledge-how, and the claim that, although success *is* an element of knowledge-how, this element is best understood dispositionally, rather than using a straight counterfactual analysis.

Finally, different forms of 'warrant' for knowledge-how will be discussed, and a distinction will emerge that is strikingly similar to that between internalism and externalism about warrant for belief. But the main thesis - the claim that counterfactual success is a necessary condition for knowledge-how - is independent of any more detailed claim about warrant. Just as internalists and externalists can agree that propositional knowledge is a special kind of true belief, whilst disagreeing about what makes it special, epistemologists should agree that knowledge-how requires a special kind of counterfactual success, before going on to discuss what makes it special.⁶

In pursuing this structural analogy between propositional knowledge and knowledge-how, we do not prejudge whether one form of knowledge is a sub-type of the other. Even though these forms of knowledge share structural features (a truth/success element, plus a warrant element), it remains open whether or not these structures are realised in different ways. For example, to be prone to non-accidental success in action may be equivalent to having a non-accidental true belief about how to act successfully, or it may at least require such a belief. The account presented here leaves open the question whether knowledge-how is distinct from propositional knowledge; instead, it provides a framework for investigating knowledge-how

directly, and thereby investigating the relationship between knowledge-how and knowledge-that.

2. Success

The relationship between knowledge-how and successful action is *prima facie* puzzling. On the one hand, knowledge-how is practical knowledge, about how to do stuff. On the other hand, it is received wisdom that someone can know how to perform without being able to perform ('those who can't do, teach'). Moreover, we know how to do things which we've never actually attempted. For example, if you've paid due attention to the safety demonstrations on planes, then you know how to put on a lifejacket, even if you've never tried it. There are, of course, many skills which one cannot acquire without hours or even years of actual practice - driving a car, tap-dancing, heart surgery. But even where repeated performance seems to be a practical necessity, performance doesn't seem to constitute knowledge-how. We would be sceptical if someone claimed both that he knew how to perform heart surgery and that he had never actually performed an operation. But that's because we know that surgery is a difficult skill to acquire, not because the impostor has contradicted himself.

Actual success is not a necessary condition for knowledge-how, but success under certain counterfactual circumstances is indeed necessary. At the limit, nobody can know how to perform an impossible task like building a perpetual motion machine - knowledge-how requires that success be possible under at least some circumstances. But which circumstances? Sarah may know how to drive even though she would not succeed if called upon to drive across a frozen lake, blindfolded and under machine-gun fire. On the other

hand, Sarah's driving a car whilst an instructor repeatedly intervenes using the dual-control pedals does not seem to qualify her as knowing how to drive. To say that Sarah knows how to drive is to attribute success to Sarah under some but not all counterfactual circumstances.

In order to understand the connection between knowledge-how and counterfactual success, it will prove useful to think in terms of tasks: there are many different driving tasks, and Sarah can know how to perform some of these tasks without knowing how to perform others. Let us look more closely at tasks and how they are structured, before proceeding with the main business of the paper.

First, different tasks correspond to different circumstances, where circumstances include states both internal and external to the subject (compare knowing how to get upstairs with a broken leg, knowing how to get upstairs in the dark, and knowing how to get upstairs in high heels). It will, however, sometimes be convenient to talk of 'families' of tasks, where the family members involve 'doing the same thing' (e.g. driving, getting upstairs, speaking Spanish) under different circumstances. Second, circumstances qualify what is known, not when the knowledge is possessed: Sarah may know how to drive-on-an-ordinary-road, even when she is in fact skidding helplessly across a frozen lake.

Third, although some tasks in a family may be more difficult than others, and knowing how to perform some may involve knowing how to perform others, not every family is neatly ordered in this way. For example, driving on snow is not clearly more or less difficult than driving in thick fog, just different. Fourth, tasks do not have sharp identity conditions. A sharp criterion for what counts as the same task and what doesn't would require an unfeasibly

sharp distinction between what counts as the same circumstances and what does not. This difficulty does not undermine the attempt to account for knowledge-how in terms of tasks: lack of clear identity criteria for tasks means we cannot say how many ‘bits’ of knowledge-how a subject has, but nor can we count how many ‘bits’ of propositional knowledge a subject has.

Sarah knows how to drive under some but not all circumstances; she knows how to perform some driving tasks, but not others. It is nevertheless tempting to ask *but does she know how to drive?* There is, however, no unique task or range of tasks which is always invoked when we ask whether someone knows how to drive. Rather, different tasks are salient in different conversational contexts. For example, in a UK context, it would be reasonable to infer from Sarah’s knowing how to drive that she knows how to drive a manual, stick-shift car. In most US contexts, however, this would not be a reasonable inference. To take a different example, a child might be said to know how to cook if she knows how to use the stove safely, whilst we would set standards higher (have a different task in mind) for an ordinary adult’s ‘knowing how to cook’, and set them higher still when discussing a chef in training.

It is clear that the same set of skills may be described as ‘knowing how to cook’ in one context and not in another. But is this really because we are talking about different tasks (requiring different skills) in different contexts? An alternative analysis would have it that each family of tasks in fact just amounts to a single task, that can be performed in different ways, or to different standards. The difference between the original account and this alternative is merely terminological - it just seems more convenient to individuate tasks in a fine-grained way. Knowing how to use the stove safely and knowing how to throw a dinner

party involve different tasks, though each may be described as ‘knowing how to cook’ in different circumstances.

A more robust challenge comes from the suggestion that each ‘family’ in fact amounts to a single task, which must be performed to canonical standards if success is to be achieved. What varies between contexts is merely the warranted assertibility of the claim ‘S knows how to cook’ (say), whilst the truth conditions remain fixed. Perhaps it is just false that the child knows how to cook, though we dignify her basic skills by kindly saying that she does know how. Alternatively, perhaps it is true that both the child and the incompetent adult know how to cook, though it would be misleading to describe the adult in those terms (it would raise false hopes of an excellent meal).

Much of what is said above, and below, could be reconstrued without loss as an account of warranted assertibility conditions for knowledge-how claims. Nevertheless the warranted assertibility manoeuvre is unattractive.⁷ For many families of tasks, it will be difficult to identify a plausible candidate for canonical circumstances or standards. Both an automatic car and a manual seem to provide perfectly reasonable circumstances for driving, for instance, and to call for different skills. Even where there does seem to be a plausible candidate for canonical or normal circumstances, we seem able nevertheless to discuss whether S knows how to X under *different* circumstances. It would be not merely misleading but outright false to say that Sheila know how to find her way home blindfolded, even though it’s true that she knows how to find her way home under normal circumstances.

This raises the question of how conversational context interacts with what is explicitly said, in order to determine which tasks are in question when knowledge-how is attributed (a similar question arises for those who think that context helps determine warranted assertibility). A full answer to this question would require a more lengthy investigation, but a few general points can be made here. First, it seems clear that context does play some such role - as in the case of manual and automatic driving in the UK and in the US, or the standards for 'knowing how to cook'. Second, we often rely on a (perhaps inarticulable) default notion of 'normal' or 'ordinary' circumstances. Failure to perform in abnormal circumstances for that task doesn't usually count against someone's possessing knowledge-how. Third, the default presumption that normal circumstances are in question may be overridden: actual, present circumstances may be more salient, or unusual circumstances may be explicitly invoked, as when someone claims to know how to find her way home blindfolded.

It is possible to know how to perform one of the tasks in a family, without knowing how to perform other tasks in that family. So in discussing knowledge-how, the target will be claims of the form *S knows how to X under circumstances C*, bearing in mind that the circumstances may not be explicitly specified. (Even the family of task may not be explicitly specified - someone might demonstratively claim to know how to do *that*.) And the central claim is that counterfactual success is a necessary condition for knowledge-how, just as true belief is a necessary condition for propositional knowledge. That's to say, if S knows how to X under circumstances C, then if S tried under to X, under circumstances C, S would successfully X.

This claim sounds like a strong one, and it prompts objections. First, it seems that one might know how to do something without having the physical capacity to perform. Second, it seems that one might know how to do something by dint of being a reliable performer, but without 100% success. There is also the rival claim that what is necessary for knowledge-how is not counterfactual success *per se*, but some kind of disposition to succeed. These are all objections to the claim that counterfactual success is necessary for knowledge-how; once these objections have been addressed, we will see that counterfactual success is not sufficient for knowledge-how, just as true belief is not sufficient for propositional knowledge.

3. Objections to the Success Condition, with Replies

First, then, what of those who know how to X, but would not succeed in X-ing if they tried, because of physical limitations? Take, for example, the sad case of the amputee cyclist. The cyclist knows how to ride a bike, and does so daily. After an accident, her leg is amputated, and she can no longer ride her bike. Yet it seems that she still knows how to ride a bike - the accident does not cause her to forget how to ride a bike, and she may be able to teach others. The amputee cyclist looks like a counterexample to my claim that if someone knows how to X, she would succeed if she tried. She looks as if she knows how, but lacks success.⁸

The apparent counterexample dissolves when we distinguish different tasks more carefully. There are two tasks here - the ordinary task of riding a bike with two legs, and the much harder task of riding a bike with one leg. Both before and after the accident, the cyclist knows how to perform the ordinary task, and does not know how to perform the more difficult task. But both before and after the accident, the cyclist has counterfactual success in the ordinary task, and not in the more difficult task. Even afterwards, if she were to try to

ride a bike, under the circumstances of having two legs, she would succeed in riding a bike. That's to say, in the nearest possible world in which she has two legs and tries to ride a bike, she succeeds.

In summary, she knows how to ride a bike with two legs and has counterfactual success in that task both before and after the accident; at no time does she know how to ride a bike with one leg, or have counterfactual success in that task. Thus we do not have a counterexample to the claim that counterfactual success is a necessary condition for knowledge-how. The cyclist is of course more restricted after her accident: beforehand, her actual circumstances are ones in which she can make use of her knowledge how to ride a bike with two legs, whereas later her circumstances have changed, rendering that knowledge useless. 'Success' worlds are more distant after the accident than before. Nevertheless, at no time is there any task with respect to which the cyclist has knowledge-how without counterfactual success.

Why did the amputee cyclist initially seem like such a plausible counterexample? The appearance of a counterexample can arise whenever a subject's circumstances are not ordinary circumstances for a given task. When we think about counterfactual success, we tend to think of nearby possible worlds (including the actual world). It seems natural to say that after the accident the cyclist cannot ride her bike, is not able to, or wouldn't succeed if she tried: we have in mind her one-legged circumstances. In contrast, when we speak about knowledge-how, we often think of normal, usual or average circumstances for a given task. We say that the amputee knows how to ride a bike, because she knows how to ride a bike under circumstances which are ordinary for most people. It's possible to invoke the other,

more difficult task by saying that she doesn't know how to ride a bike in her present circumstances or that she doesn't know how to get herself cycling again.

The same strategy will handle apparent counterexamples regarding coaches (who know how to perform under the circumstances of being younger, fitter or more agile), prisoners (who know how to perform if free) and those who lack material resources (who know how to perform given resources). This strategy may seem too powerful: do we all know how to speak Russian, given that we would succeed given years of practice? This concern will be addressed in section 7; at present we are merely trying to establish necessary conditions for knowledge-how.

The second type of counterexample concerns those who seem to know how to X under circumstances C, even though they do not invariably succeed, even under those circumstances. Sandra is a highly competent baker, who knows how to make delicious bread. Of course, she would not succeed if no raw materials were available, or if she had her hands tied behind her back, but such counterfactual failure doesn't seem to count against her: she knows how to make delicious bread under normal bread-making circumstances. Even so, it might seem that she doesn't satisfy the counterfactual success condition for knowledge-how, because occasionally she fails to make delicious bread when she tries, even under quite ordinary circumstances. After all, nobody's perfect. There are two options here. The first is to treat these cases like that of the amputee cyclist, arguing that once we pinpoint exactly what the subject knows how to do, we will see that she invariably succeeds under the relevant circumstances. The second option is to weaken the success condition. I will examine these options in turn.

The first option is to say that Sandra is a candidate for knowing how to make delicious bread under just those circumstances in which she succeeds. She may, if she also satisfies some ‘warrant’ condition, know how to make delicious bread when ingredients are available, she has normal use of her body, she is concentrating, the oven doesn’t break down half way through baking,...there is no need to fill in the dots. This characterisation of what Sandra knows how to do is not viciously circular. This is in part because knowledge-how requires warrant as well as success. But it is also because it is common in human affairs to define competence in terms of competent performers: to be a competent performer is to succeed under circumstances under which a competent performer would succeed.⁹ Many of us would not succeed under those circumstances, for we are not competent bread-makers.

The second option is to weaken the counterfactual success condition, arguing that if S knows how to X under C, then S would *usually* succeed in X-ing were S to try under circumstances C. Presumably, there is no exact threshold which qualifies the subject as knowing how to X under C.¹⁰ Again, the threshold may be set by reference to competent performers. Is the threshold ever very low - could someone know how to do something even though she succeeded only rarely under the relevant circumstances? It seems unlikely, although of course the way is open for possible counterexamples. We can handle occasional success either by describing the subject as knowing how to perform only under restricted circumstances, or else by saying that she does not know how, despite her occasional success.

We have seen two ways in which know-how-with-occasional-failure may fit into the account - either we can ‘fix’ the circumstances, so that success is always obtained, or else we can

argue that knowledge-how requires reliable but not exceptionless success. The first option has the advantages of theoretical simplicity, retaining a straight counterfactual success condition, whilst allowing some leeway in which task is discussed. The second option introduces some fuzziness into the standards for knowledge-how, rather than into what is known. But this may be unavoidable (as it is with propositional knowledge) since the non-accidentality, warrant or justification which is required to make the difference between success and knowledge may itself allow of borderline cases. The reader may choose between these two options; the choice will not make a difference to what follows. Counterfactual success (whether 100% or less) in performing a task is a necessary condition for knowing how to perform that task.

4. Dispositions to Succeed

The present claim is that if S knows how to X under C, then if S tried to X, under C, S would succeed. A different view is that if S knows how to X under C, then S is disposed to succeed if she tries to X under C. These look like different claims about what stands to knowledge-how as true belief stands to propositional knowledge. But do the accounts differ significantly? A disposition is actually possessed by an object, even when unmanifested. But, similarly, a counterfactual is actually true of an object, even when the antecedent is not satisfied: it is actually true right now that if Stella were to try to make an omelette, she would succeed.

Another thought is that an object's possessing a certain disposition may be explained by its possessing a grounding property - an object's fragility may be explained by its molecular make-up (may even be identical with that make-up). But, similarly, we may explain why

Stella would successfully make an omelette if she tried, by mentioning her mental states, her physical properties, and those of her surroundings. Relatedly, we sometimes explain counterfactual success by alluding to knowledge-how: she would succeed if she tried because she knows how to do it. But such explanations are not ruled out by the claim that counterfactual success is a necessary condition of knowledge-how: knowledge-how is not thereby reduced to success, and counterfactual success may be explained in terms of actual properties. Moreover, we can sometimes explain individual success by fitting it into a pattern - Stella would make an omelette right now if she tried, because in general if she tries to make an omelette she succeeds.

Finally, a dispositional account might seem attractive because knowledge-how seems liable to 'finkishness', just like fragility and solubility. A disposition is finkish iff it would be absent under manifestation conditions (Lewis 1997). For example, a vase might be fragile, but caused to be non-fragile by being dropped: though it is fragile, if it were dropped it would not break. It looks as if knowledge-how may be finkish: Sylvia knows how to get to her home from the city centre, but if she were in the city centre she would have a panic attack and forget how to get home. If she were to try to get home from the city centre she would not succeed.

It is uncontroversial that ordinary dispositions may be finkish and thus that straight counterfactual analyses of dispositions fail. But we need not accept that knowledge-how may be finkish. Like the amputee cyclist, Sylvia's knowledge-how matches her counterfactual success. Sylvia does not know how to get home from the city centre under the circumstances of being prone to panic attacks. She does know how to get home from the city centre under

circumstances which are normal for most people, but she also satisfies the counterfactual success condition for this task: if she were to try under such circumstances, she would succeed in getting home.

The dispositional success condition preserves the spirit of the preferred account - the connection between success and knowledge-how - though it differs in the letter. But replacing the straight counterfactual success condition with a dispositional success condition does not seem to offer any advantages, and it raises additional questions about the nature of dispositions. Let us therefore set this dispositional suggestion aside.

5. Success isn't Everything

Counterfactual success, though necessary, is not sufficient for knowledge-how. Such success may be accidental, or somehow ill-grounded, much as the truth of a belief may be accidental or the belief ill-grounded. Sections (5) and (6) will establish that not all success amounts to knowledge-how, then section (7) will consider the harder question of what it takes for success to be appropriately non-accidental.

First, success cannot amount to knowledge-how unless intentional action is involved. We don't describe ourselves as knowing how to produce white blood cells. At a push, we might say that someone knows how to produce white blood cells if she knows what sort of diet or drugs might promote this process, but to this extent producing white blood cells is indeed an intentional action. Just as propositional knowledge requires representation, knowledge-how requires intentional action. Fortunately, the success condition explains this: if S knows how,

S would succeed if S tried. 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.

The ‘trying’ clause also explains the opacity of knowledge-how attributions: to flip the switch is to alert the prowler, so if Sarah is able to do one of these, she is able to do the other (Davidson 1963). Yet it seems that Sarah could know how to turn on the light without knowing how to alert the prowler, if it doesn’t occur to her that turning on the light will alert the prowler. Here, Sarah satisfies the success condition for turning on the light, but not for alerting the prowler, and this explains the difference in her knowledge-how.¹¹

Second, it is easy to think of exotic tasks where each of us has counterfactual success. For example, Shula satisfies the success condition for the tasks of speaking-Russian-after-five-years-of-living-in-Moscow, flying-a-jet-after-much-training and the like. If she tried to speak Russian or fly a jet in such distant circumstances, she would succeed. Yet the mere fact that Shula *could learn* how surely doesn’t mean that she now knows how to speak Russian or how to fly a jet. Thus far, Shula does not present a problem: she does not have counterfactual success in the task of speaking Russian under normal circumstances, at will. (These circumstances are normal in that performance at will would be expected from someone described as knowing how to speak Russian.) Yet she does have counterfactual success at the more qualified task, and some may be reluctant to ascribe knowledge-how to her even there. After all, Shula could satisfy the success condition for speaking-Russian-after-five-years-in-Moscow even if she has never heard of Russian or Moscow. Recall, however, that in addition to counterfactual success, some kind of warrant is also required for knowledge-how. It may be that if she has not heard of Russian, Shula cannot satisfy the warrant condition for

knowing how to speak Russian after five years in Moscow, even if she satisfies the counterfactual success condition; warrant will be discussed in section (7).

Third, someone might satisfy the success condition for driving (if she tried, she'd succeed) without knowing how to drive, just because she would try to drive *only if* she would succeed.¹² As with the Russian case, two compatible strategies are available to me. One is to suggest that this person may satisfy the counterfactual success condition, but not the warrant condition. The other is to distinguish tasks more carefully. This person does not satisfy the success condition for driving-without-further-training; the apparent difficulty arises because the nearest world in which she tries to drive without further training (where she fails) is more distant than the nearest possible world in which she tries to drive (where she succeeds). She neither knows how to drive under normal conditions (at will, without further training) nor has counterfactual success in driving under these conditions. She does have counterfactual success in driving after further training, and some might be reluctant to attribute know how even in this task; as with the Russian case, such people will need to invoke her lack of warrant.

6. Accidental Success

Just as it is possible to have an accidentally true belief without thereby having propositional knowledge, it is possible to have counterfactual success that is in some sense 'accidental', without thereby having knowledge-how. This section describes three cases of accidental success, then asks what is missing in such cases, what prevents the subjects involved from qualifying as knowing how. The three cases have rather different features. The reader need

not be convinced by all three: one is enough to make the point that success does not always suffice for knowledge-how.

First, consider Sally, out for an ill-advised winter walk in the hills. She has no idea what to do in the event of an avalanche. When an avalanche occurs, she mistakes the snow for water, makes swimming motions, and - luckily - escapes the avalanche, since in fact the way to escape an avalanche is to make swimming motions. Sally satisfies the counterfactual success condition for escaping avalanches - if she were to try, she would succeed, for she is prone to mistaking snow for water. Yet it seems that her success is merely a matter of luck: her two failings (not knowing how to escape from an avalanche, mistaking snow for water) cancel each other out. Despite her success, she does not know how to escape avalanches.

Second, 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 doesn't know how to annoy Joe, perhaps because she misconstrues the situation.

Finally, consider Shelley, who attempts to make a cake by throwing together the ingredients she discovers on opening the kitchen cabinet. Luckily for her, the cupboard contains flour, sugar, butter and eggs, and she makes a passable cake. Yet it seems that Shelley does not know how to make a passable cake. Of course, she doesn't know how to make a cake under normal circumstances, which include selecting the ingredients. But nor is she like a certain blind person, say, who knows how to make a cake under the restricted circumstances of being provided with the right ingredients. Shelley doesn't even know how to do that, even though

she would succeed if she tried under such circumstances. If she notices her success, and recalls her method, she may come to know how to make a cake. But the mere fact of her lucky success does not initially qualify her as knowing how to make a cake.

Of course, for some tasks, accidental success is practically impossible: nobody could simply guess how to operate a space shuttle. And for some propositions, accidentally true belief is practically impossible: nobody could simply guess all the telephone numbers in Mexico City. But these practical limitations do not show either that propositional knowledge just is true belief, or that knowledge-how just is counterfactual success.

7. Knowledge-How and What it Takes

Not all true belief amounts to propositional knowledge, and not all counterfactual success amounts to knowledge-how. In each case, success must be in some sense non-accidental, or warranted, where ‘warrant’ is a neutral term for whatever makes the difference between true belief and knowledge, or between counterfactual success and knowledge. In many cases of accidental success without knowledge-how, there are two broad respects in which the subject is deficient. First, she lacks understanding of her own success. Sally falsely believes that the snow is water, Susie is mistaken about what annoys Joe, and Shelley is simply guessing how to make a cake. Second, success in these cases seems fortuitous or accidental. For example, the subject does not act as she does *because* that’s a good way of achieving her goal, each would have behaved in a similar way even had that not been a good way of achieving their goal, and each obtained her method from a source (guesswork, misconstrual of the situation) which is an unreliable source of successful methods.

Success may fall short of knowledge-how in two ways: through a subject's failure to understand her situation, or through the absence of a suitable connection between the goodness of the method and the subject's use of that method.¹³ These categories of deficiency are suggestively similar to ways in which 'mere' true belief may fall short of propositional knowledge: through a subject's failure to grasp good reasons for her belief, or through the absence of a suitable connection between the truth of the belief and the subject's holding that belief. This is a powerful analogy: the various internalist and externalist accounts of propositional knowledge provide raw material for analogous accounts of knowledge-how.¹⁴

Do the two categories of deficiency correspond to rival accounts of warrant for knowledge-how, or to compatible accounts? We need to look more closely at the possibilities. First, the 'internalist' requirement that the subject have some understanding of her success or actions must not be implausibly strong. As Ryle (1949) argued, one can know how to do something without having detailed knowledge about one's method. A more modest requirement is that the subject know she would succeed if she tried. Such knowledge might be underpinned by an understanding of method, or it might be gained through straightforward self-observation (whatever I'm doing, it seems to work). We hesitate to attribute knowledge-how to the annoying smoker who misunderstands her own actions, even if she has observational evidence of her success: perhaps her false beliefs about how she succeeds undermine her inductive justification for her true belief that she will succeed. Is the internalist requirement still too strong? Couldn't someone know how to do something even if she'd never stopped to reflect on her success? If so, then we should require merely that the subject be warranted in believing that she'd succeed, whether or not she has actually formed such a belief.

How should we spell out the ‘externalist’ requirement that there be a suitable connection between the method’s being a successful one, and the subject’s using that method? There are many possible accounts of what makes a connection ‘suitable’, but there is also a more general question. Must the connection hold in the world in which the subject knows how, or in the world where the subject uses her method? For example, suppose we require a causal connection from the goodness of the method to the use of the method. This may be spelt out in two different ways: either we require that if the subject had used the method, that would have been because the method was a good one, or else we require that the subject’s actual tendency to use the method be caused by the method’s actual tendency to be a good one. Often these requirements will be met together: the actual cause of someone’s actual readiness to act in a certain way has a counterpart which is a counterfactual cause of her counterfactual action. But whether these requirements always coincide will depend upon the nature of the ‘suitable’ connection, upon how counterfactuals are actually grounded, and upon exactly which counterfactuals are in question: careful specification of tasks solves related problems concerning finkishness.

What, then, is the relationship between the internalist and the externalist requirements? If someone is entitled to believe that she would succeed if she tried, does this guarantee a suitable connection between the goodness of the method and the subject’s use of the method, either in this world or in the world where she tries? This may depend upon what counts as a ‘suitable’ connection, but it seems unlikely that satisfaction of the internalist requirement could guarantee satisfaction of the externalist requirement. It is often the case that a subject’s knowing that her method is a good one explains the fact that she would use it. But someone

could be entitled to believe or even know that her method is a good one without this having any causal, explanatory or counterfactual link to the fact that she would use it. She might for some reason be compelled to use that method. So satisfaction of the internalist requirement will often but not always involve satisfaction of an externalist requirement.

The link seems even weaker in the opposite direction, although again the details will depend upon what is required of a 'suitable' connection. Even if a subject's tendency to use a method is suitably connected to its being a good method, the subject will not in general be entitled to believe that her method is a good one, for we do not in general have privileged access to the explanations of our actions. The internalist and externalist requirements on knowledge-how are distinct, yet both are *prima facie* plausible - there is some deficiency in someone who fails to understand her situation, and there is some deficiency in someone whose use of a method is not connected to the fact that it is a good means to her end. But both types of account need further development - for the moment it remains unclear whether they correspond to two related concepts of knowledge-how, or whether they are rival accounts of a single concept, and if so, which we should prefer.

Now that we have begun to investigate warrant for knowledge-how, we can look again at tasks like speaking-Russian-after-five-years-in-Moscow. It is easy to have counterfactual success in such tasks; how easy is it to have warrant? Subject, of course, to the details of different accounts, it looks as if externalist requirements may be satisfied. In the world in which Shula has lived five years in Moscow, and consequently speaks Russian, it is no accident that she does so. What of the internalist requirement? Is Shula entitled to believe that she would successfully speak Russian if she tried after five years in Moscow? On the

one hand, she knows that she has average language-learning abilities, and that most people pick up a foreign language if they are immersed in it. On the other hand, if she has never heard of Russian, this may mean she is not entitled to *any* beliefs about Russian, and thus does not know how to speak-Russian-after-five-years-in-Moscow. It seems plausible that, if Shula has heard of Russian, and knows of her language-learning skills, then she does know how to speak-Russian-after-five-years-in-Moscow; it's just that this is widespread and relatively useless knowledge-how, and thus not usually worth mentioning. Those who wish to insist that Shula does not have know how with respect to this task, even if she has heard of Russian, will need to place fairly stringent restrictions on the internalist requirement for knowledge-how; perhaps Shula needs a fuller understanding of what is involved in speaking Russian.

We have seen, then, that knowledge-how does require counterfactual success in action, that knowledge-how shares general structural features with propositional knowledge, and that the possible accounts of warrant for the two categories of knowledge have some striking similarities. This indicates several things: first, that our notion of knowledge combines ideas of success with ideas of non-accidentality and that it would be worth investigating 'acquaintance' knowledge with this in mind; second, that we cannot fully understand knowledge-how without looking more closely at the various different accounts of warrant; and third, that where knowledge-how is invoked outside epistemology, more caution is required until we are clearer about warrant and what it requires.

Bibliography

- Austin, J.L. (1956): 'Ifs and Cans', reprinted in his *Philosophical Papers*, Oxford: Oxford University Press, 2nd edition 1970, pp. 205-232.
- Blackburn, Simon (1996): 'Securing the Nots: Moral Epistemology for the Quasi-Realist', in Sinnott-Armstrong and Timmons (eds.) (1996), pp. 82-100.
- Carr, David (1979): 'The Logic of Knowing How and Ability', *Mind*, 88, pp. 394-409.
- Collins, H.M. (1985): *Changing Order*, London: Sage.
- Craig, Edward (1990): *Knowledge and the State of Nature*, Oxford: Oxford University Press.
- Davidson, Donald (1963): 'Actions, Reasons and Causes', *Journal of Philosophy*, 60.
- DeRose, Keith (1998): 'Contextualism, an Explanation and Defense', in Greco and Sosa (eds.) *The Blackwell Guide to Epistemology*, Oxford: Blackwell (1998).
- Dummett, Michael (1991): *The Logical Basis of Metaphysics*, London: Duckworth.
- Ginet, Carl (1975): *Knowledge, Perception and Memory*, Boston: Reidel.
- Goldman, Alvin (1976): 'What is Justified Belief?', in G.S. Pappas (ed.) *Justification and Knowledge*, Dordrecht: Reidel.
- Lewis, David (1988): 'What Experience Teaches', *Proceedings of the Russellian Society*, 13, pp. 29-57.
- Lewis, David (1997): 'Finkish Dispositions', *The Philosophical Quarterly*, 47, pp. 143-158.
- Mellor, D.H. (1993): 'Nothing Like Experience', *Proceedings of the Aristotelian Society*, XCIII, pp. 1-16.
- Moore, A.W. (1997): *Points of View*, Oxford: Oxford University Press.
- Plantinga, Alvin (1993): *Warrant: The Current Debate*, Oxford: Oxford University Press.
- Ryle, Gilbert (1949): *The Concept of Mind*, London: Hutchinson.

Sayre-McCord, Geoffrey (1996): 'Coherentist Epistemology and Moral Theory', in Sinnott-Armstrong and Timmons (eds.) (1996), pp. 137-189.

Sinnott-Armstrong, Walter and Timmons, Mark (eds.) (1996): *Moral Knowledge?*, Oxford: Oxford University Press.

Stanley, Jason and Williamson Timothy (2001): 'Knowing How', *Journal of Philosophy*, XCVIII pp. 411-44.

Williamson, Timothy (2000): *Knowledge and its Limits*, Oxford: Oxford University Press.

¹ I have presented papers based on this material at the universities of Aberdeen, Dundee, Durham, Edinburgh, Glasgow, Leeds, Lund, Michigan, Reading, Singapore, St Andrews, Stirling, and Warsaw, and I thoroughly appreciate the cross-examinations I received on those occasions. Many people have also helped me through discussion or written comments; I am particularly grateful to Edward Craig, Sam Inglis, Peter Kail, Martin Kusch, Philip Lakelin, Peter Lipton, Ian Macho and Hugh Mellor.

² Recent exceptions include Craig (1990), Moore (1997), and Stanley and Williamson (2001).

³ For examples, see Dummett (1991, chapter 4), Blackburn (1996), Sayre-McCord (1996), Lewis (1988), and Mellor (1993).

⁴ Following Plantinga (1993), 'warrant' is used here as a *neutral* term for whatever makes the difference between mere true belief and propositional knowledge, or between mere success and knowledge how. It remains open at this stage whether the same features can play the role of warrant in the two cases, whether 'warrant' is reductively analysable in either case (Williamson 2000), and what relationship there is between warrant and justification.

⁵ Counterfactual success in action is here taken to play a role in knowledge-how analogous to

that played by true belief in propositional knowledge. A stronger claim would be that believing truly is a form of successful action, or even that believing truly is a matter of counterfactual success in occurrent belief. While these further claims are attractive, they are not presupposed in this paper.

⁶ It is plausible to identity propositional knowledge with some special kind of true belief, but is it plausible to *identity* knowledge-how with a special kind of counterfactual success, even if having that kind of success is both necessary and sufficient for knowledge-how? This paper does not provide an argument for that final step; indeed the matter may be resolved differently according to different accounts of warrant.

⁷ In defending contextualism about propositional knowledge, Keith DeRose (1998) discusses various warranted assertibility manoeuvres. The present contextualism about knowledge how differs from standard contextualism about propositional knowledge: here, what is said to be known varies with the context (successful action plays the role of true belief). It is a further question whether the standards for knowing/warrant *also* vary with the context.

⁸ Stanley and Williamson (2001, 416), following Ginet (1975), take cases like this to show that “ascriptions of knowledge-how do not even entail ascriptions of the corresponding abilities”. As Austin (1956) demonstrates, the semantics become even more complex and context-sensitive when we talk of what a subject *can* do or is *able* to do. Discussion of counterfactual success is intended to avoid these difficulties.

⁹ The ‘experimenters’ regress’ reflects this: significant experimental data – as opposed to artefacts and errors – are those produced by competent experimenters using appropriate equipment, but the significance of data and the competence of experimenters cannot be assessed entirely independently. Collins (1985).

¹⁰ This echoes reliabilism: a reliable process is one which *usually* produces true beliefs.

Goldman (1976).

¹¹ Carr (1979) discusses the opacity of knowing-how attributions in more detail.

¹² This raises the possibility of finkishly absent know-how. Introducing dispositions won't help: she's disposed to successfully drive if she tries.

¹³ An action may be 'warranted' in a different sense - worth doing - even if its success is accidental, unearned or 'unwarranted' in my epistemic sense. If the penalties for inaction are high, it can be worth guessing how to do something, and a lucky guess may lead to accidental success.