King's Research Portal

DOI:

[10.1111/1467-9752.12256](https://doi.org/10.1111/1467-9752.12256)

Document Version

Peer reviewed version

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Lum, G. (2017). Making Sense of KnowingHow and KnowingThat. *JOURNAL OF PHILOSOPHY OF EDUCATION*, 51(3), 655-672. <https://doi.org/10.1111/1467-9752.12256>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Making sense of knowing how and knowing that

Gerard Lum

Noting the resurgence of interest in Ryle's knowing how / knowing that distinction the paper examines the viability of this distinction along with claims to reduce one form of knowing to the other. It is suggested that by separating questions about the nature of knowledge from questions about third person attributions of knowledge, it is possible to determine the circumstances in which use of the distinction is and is not viable. By drawing attention to one important feature of Ryle's argument and suggesting a corrective to another it is argued that by Ryle's own account the distinction cannot properly be regarded as an epistemological distinction. The paper concludes by considering some negative educational implications of the assumption that it is possible to distinguish two kinds of knowledge.

The last decade or so has seen a resurgence of interest in Ryle's knowing how / knowing that (KH/KT) distinction, prompted in no small part by Stanley and Williamson's (2001) provocative intellectualist reading of the distinction. Recent work has seen the battle lines drawn between those who essentially defend Ryle's original distinction, or some variant of it, and those who would claim that one form of knowing reduces to the other. It is not my intention here to rehearse the arguments of the various protagonists but rather to examine afresh the viability of the idea that knowledge is of essentially two kinds. By drawing attention to one particularly crucial feature of Ryle's argument – well known yet conspicuously absent from much of the KH/KT debate – and suggesting a corrective to another, I will argue that even by Ryle's own account the distinction cannot properly be regarded as an *epistemological* distinction, that is, as demarcating two different kinds of knowledge. However, this is not to say that we are necessarily mistaken in our ordinary use of the terms knowing how and knowing that. Indeed, it turns out to be of no small significance that the distinction is deeply rooted in our language and our attempts to articulate what it is that other people know. What matters, I will argue, is being clear about where our use of the KH/KT distinction *does* make sense and where it *doesn't*.

REVISITING RYLE'S USE OF THE DISTINCTION

Ryle's use of the KH/KT distinction in *The Concept of Mind* is clearly bound up with his wider ambitions. Paul Snowden (2004) has noted how Ryle's use of the distinction is closely connected with his attempt to refute the 'Intellectualist Legend', the idea that 'the intelligent execution of an operation must embody two processes, one of doing, and another of

theorising' (Ryle, 1949, p. 32). And Stanley and Williamson (2001) have similarly focused on Ryle's efforts to dismiss the intellectualist legend. Yet it is important not to lose sight of Ryle's still broader purpose, that of demonstrating the alleged misuse of a single concept, along with the range mental epithets that people use in the supposedly mistaken belief that in so doing they are referring to some inner mental realm, some ostensible 'ghost in the machine' (*ibid.*). The difference between these differently construed aims is crucial because they suggest two fundamentally different kinds of question. Indeed, one difficulty with the *The Concept of Mind* is precisely its ambiguity, that is, whether Ryle's overarching concern is with the nature of mind and knowledge, or with our *descriptions* of particular minds, that is, with our use of mental epithets. Even Ryle's title is ambiguous on this score. But what *is* clear in Ryle, and is stated as the express purpose of his chapter on knowing how and knowing that, is his intention 'to show that when we describe people as exercising qualities of mind, we are not referring to occult episodes of which their overt acts and utterances are effects; we are referring to those overt acts and utterances themselves' (p. 25). This aspect of Ryle's argument is certainly well known, yet its implications for the KH/KT distinction seem not to have had the attention they deserve. For if in describing people as 'knowing how' or 'knowing that' we are in fact referring to 'overt acts and utterances' rather than knowledgeable states of mind then there may be an important sense in which by Ryle's own account these terms cannot be assumed to represent an *epistemological* distinction.

It is clear from the passage where Ryle introduces the distinction that the terms knowing how and knowing that are not meant to have any special technical meaning; rather, these are simply terms used in ordinary language. As he says, 'We speak of learning how ... as well as learning that...; of finding out how as well as finding out that' (p. 28). Ryle's interest here is with 'our *descriptions* of people as knowing' (my emphasis; *ibid.*), with how we *speak* about what other people know. He sees this distinction as connected with the distinction that can be made between people's 'operations' and 'the truths that they learn' (*ibid.*); in other words, we speak of a person's knowing how or knowing that according to whether what they know is made manifest in the form of acts or in the form of utterances. Ryle does not venture to say why we should have an interest in making this distinction, but it seems reasonable to suggest that by thus indicating *how* a person's knowledge is made manifest we are able to distinguish more clearly between a person's knowing one thing and knowing another. It allows us to distinguish between, say, a case of knowing how to ride a bicycle and a case of knowing facts about cycling in a way that that would not be clear if we were simply to talk of

a person having ‘knowledge of cycling’ (see Lum, 2009). Yet not all manifestations of knowing are amenable to being categorised in this way: the writing of an essay, for example, being a case in point. Such examples can be explained by recognising their ambiguity *as manifestations*, being at once both ‘act’ and ‘utterance’, to use Ryle’s terminology. For the moment suffice it to say that we make a fatal move if we leap too quickly to the idea that such cases demonstrate the application of two kinds of knowledge.

Certainly there is cause to be circumspect about taking commonplace references to knowing how and knowing that too literally. If someone says ‘John knows that the key is under the mat’ in response to the question ‘Does John know how to get in?’ they would be attributing to John rather more than knowledge of a fact; they would in effect be saying that ‘John knows how to get in’. Similarly, we might say of someone that they know how to answer a question or know how to say the right thing when we mean nothing other than they know *that* the answer is such and such, or know *that* they should say certain things. But such matters aside, the main point here is that we can acknowledge that the KH/KT distinction does have perfectly viable applications, in English at least, when used to indicate the manner in which a person’s knowing is made manifest and thus indicate more clearly what it is they know.

Now Ryle’s point as regards our lack of access to other minds is incontrovertible. The difficulty, of course, as Ryle was all too aware, is that it seems tantamount to behaviourism. More specifically, it leaves us incapable of explaining how we are able to distinguish cases perhaps identical in terms of overt behaviour yet seemingly very different in terms of what is known. To use Ryle’s example, how are we able to distinguish the skilful trippings and tumblings of the clever clown from the ‘visibly similar trippings and tumblings of the clumsy man’ (p. 33)? As is well known, Ryle’s solution was to invoke the notion of dispositions: what distinguishes the clown is his disposition to act the way he does. The problem here lies with Ryle’s explanation of ‘disposition’:

To possess a dispositional property is not to be in a particular state, or to undergo a particular change, it is to be bound or liable to be in a particular state, or to undergo a particular change, when a particular change is realized. (Ryle, p. 43)

The difficulty with this – despite Ryle’s insistence that to have a dispositional property is ‘not to be in a particular state’ – is precisely that it encourages the thought that what is at issue is some feature or ‘property’ of the knower. This is where the question of Ryle’s substantive

agenda becomes crucial. If Ryle is trying to say something about the nature of mind then it might be apt for him to talk about the knower possessing certain properties. Yet this is not what Ryle has led us to believe; the thrust of his argument thus far has centred not so much on the nature of mind but, rather, on our *descriptions* of particular minds, what it is possible for us to mean when we apply intelligence epithets, and so on.

The important thing here is that questions about the nature of mind and knowledge are logically distinct from questions about what it is possible to know and say about the mind of particular person or about what a particular person knows. In much of the discussion surrounding this topic the difference between these two kinds of question tends to be overlooked, indeed Ryle himself can be seen continually to conflate these two kinds of question. But I want to suggest that in order to be clear about the KH/KT distinction, where the distinction is and is not valid, it is vital to separate these two sorts of question. In education the difference has an obvious practical bearing, the one kind of question being pertinent in connection with the curriculum, with how we conceive of what learners need to know, the other being relevant in connection with assessment and the processes by which it is possible to determine what particular learners know. And we can take Ryle's puzzle of the clown/clumsy man as posing a question of the latter kind, being essentially a question about how it is possible, given our patent lack of access to other minds, to determine what a person knows in contradistinction to how they behave.

Now there is an important sense in which the facility to determine what a person knows in contradistinction to how they behave is fundamental to the business of teaching. Consider the not at all unusual situation in which a teacher, on asking a pupil a question and receiving the correct answer, just *knows* that he doesn't really know. Or conversely, she might feel confident that the pupil *does* know in spite of *not* receiving the correct answer. If asked whether or not the pupil knows she would rightly say that it depends on what we mean by 'know'. She would likely explain the matter in quasi ontological terms, as a discrepancy between the 'outer' and the 'inner', as a difference between the pupil's outward behaviour and what he 'really' knows. But of course this cannot be right. In this much Ryle is surely correct: the teacher has no privileged access to the realm of the mental, the only evidence she has is that of outward behaviour.

The idea that the pupil might possess a certain disposition or complex of dispositions seems on the face of it to be of little help here. How might such a disposition be discerned? Perhaps, we might suppose, this particular instance of behaviour is the latest in a series of similar instances and the teacher has made some sort of inference or calculation based on those past cases in order to surmise that the pupil has a certain ‘dispositional property’. Yet this is unsatisfactory for a number of reasons.¹ Not least, it does nothing to divest us of the apparent need to distinguish between the ‘outer’ and the ‘inner’; it would still be necessary to distinguish between the pupil’s behaviour and what he ‘really’ knows in order that such cases could be regarded as ‘similar instances’. And this kind of explanation fails to account for situations in which there is no history of similar cases. As it stands, Ryle’s appeal to the notion of dispositions appears lacking and if we are going to be clear about the ways in which the KH/KT distinction may or may not be valid we need a more plausible explanation. Accordingly, I want to suggest a corrective to Ryle’s scheme of things which might allow us to partially rehabilitate his notion of dispositions, although I should stress that it is far from certain that Ryle would be in agreement with what I am going to propose.

A CORRECTIVE TO RYLE’S NOTION OF DISPOSITIONS

If I am asked whether my elderly aunt would be able to negotiate the London underground so as to get from Paddington Station to Russell Square, something she has never done before, my answer would be based not on any calculation of previous cases, but rather, on my *understanding* of my aunt and her capabilities. Donald Davidson (2001) has commented on the extraordinary facility we have to draw on any diverse and fragmentary evidence we have about a person and somehow ‘assemble such material into a convincing picture of a mind’ (p. 15). It would seem that the processes involved in our doing this are largely tacit and unconscious; as Davidson says, ‘we know how to do it without necessarily knowing how we do it’ (*ibid.*). Indeed, there is ample empirical evidence in the psychology of perception to suggest that this is something we instinctively and automatically do for each and every person we know.²

I have suggested elsewhere (Lum, 2012) that this has unrecognised implications for educational assessment, not least because it suggests that whenever we set out to determine what a person knows we have at our disposal not one but *two* logically distinct kinds of judgement. On the one hand we might bring to bear *judgements of identity* whereby we seek merely to confirm the presence or otherwise of some specific, predetermined behavioural

manifestation (e.g. 'he gave the correct answer'). This is the way we ordinarily think of assessment, particularly when employing formal processes centred on 'objective' criteria.³ But in addition to this we can also be seen to employ what I have dubbed *judgements of significance* by which we actively draw on *any* available evidence that we judge to be relevant to the matters in hand. It would seem that by this means we are able, in effect, to construct a 'picture of a mind' which, in turn, allows us to judge what a person knows (e.g. 'but he didn't really know'). Such 'pictures' are constantly updated as evidence becomes available, with judgements of identity continually being monitored as to whether they are consistent or at variance with the 'picture' we have of a person and their capabilities. Since by this means we are able to judge what a person may or may not be disposed to do in particular circumstances we can partially reinstate Ryle's notion of dispositions, not as a property of the person but, rather, as part of the 'picture' we have of the person.

Whether Ryle would accept this modification is, as I say, far from certain. The essential difficulty for Ryle is that his positivistic leanings often cause him to focus on the supposed 'object' of attention, i.e. 'the knower', when the focus should more properly be on the *observer* and the processes brought to bear by the observer in making judgements about what a person knows.⁴ That said, Ryle does concede something of an interpretative role for the observer in acknowledging that in order to understand a person's 'deeds and words' it is necessary to have 'some degree of competence' (Ryle, 1949, p. 54) in common with the performer. What I am proposing is certainly consistent with Ryle's thesis about our lack of access to other minds, whilst being such as to avoid the accusation of behaviourism. It also explains the folk psychology of mental epithets, our apparent conviction that we can know and say things about a person's mind in contrast to their behaviour. What is missing in Ryle is the sheer extent of what the observer brings to the situation: in the case of the clown/clumsy man an understanding of the kind of things that happen in circus tents, certain expectations as to how people dressed in certain costumes are likely to behave, and so on. That this is so is evidenced by the possibility that the observer can get it wrong: if the clown were to have a nasty unintended fall the audience may well see it as part of the act; were he to dress in a business suit and perform his routine in the high street it is likely that concerned bystanders would come to his aid. But the more immediate thing is that our acknowledging these two kinds of judgement allows a more plausible explanation of how it is possible to determine what a person knows in contradistinction to what they say or do. And this in turn,

I want to suggest, allows for a more coherent account of the notions of knowing how and knowing that in relation to attributions of knowledge.

A QUESTION OF KNOWLEDGE ATTRIBUTION

We can concur with Ryle when he says, in effect, that our judgements about what other people know must necessarily be based on manifest evidence. What matters – and what is conspicuously missing in Ryle – is what we *do* with that evidence, whether we confine our judgements to *specific* behavioural manifestations, or whether our judgements are such as to draw on *any* relevant evidence by way of constructing a ‘picture of a mind’. It is precisely because Ryle appears to confine attention to specific ‘acts and utterances’ that he has been taken to task for seeming to equate knowledge *with* those acts and utterances.⁵ And this has generally come to be seen as the main bone of contention between the anti-intellectualist and intellectualist positions, with those who hold the latter position being at pains to show that outward behaviour is either not necessary or not sufficient for knowledge. Indeed, it might be said that much of the burgeoning debate surrounding the KH/KT distinction over the past decade or so has been less about the relation between these two ostensible forms of knowledge than about the relation between know-how and ability,⁶ with particular puzzlement being caused by the fact that attributions of know-how sometimes entail ability attributions and sometimes do not.⁷

Stanley and Williamson offer the following much discussed illustration in their attempt to distance know-how from ability:

...a ski instructor may know how to perform a certain complex stunt, without being able to perform it herself. Similarly, a master pianist who loses both of her arms in a tragic car accident still knows how to play the piano. But she has lost her ability to do so.

(Stanley and Williamson, 2001, p. 416)

The suggestion, of course, is that in each case we would be bound to concede that the person has knowledge even though they are not able to do the thing in question. Yet this is misleading. First, we need to be clear that this is a question about the *attribution* of knowledge; that is, a question of whether and how we would be justified in saying that knowledge obtains in a particular case. Seen thus we can recognise that the two examples are actually very different. How would we know that the ski instructor ‘know(s) how to perform a certain complex stunt’? Presumably, we might suppose, because she can evidently teach her students to perform the stunt. Accordingly, this is *not* a case of knowing without ability

but, rather, a case in which it is necessary to acknowledge the difference between knowing how to do one thing and knowing how to do another.⁸ What the ski instructor knows is how to instruct her students to do the stunt, and this knowledge is different from the knowledge required to perform the stunt. In a different case the reverse might be true: someone might be able to do the stunt and yet not know how to teach others how to do it. Again, it would be a case of knowing one thing as opposed to knowing something else. As is so often the case in the literature on this topic, the matter comes down to ambiguities of language, in this case the deceptive ambiguity of 'know how to perform'.⁹

The case of the pianist is different. But here again, with our modification in place, Ryle can be vindicated. On what grounds would we judge that the master pianist knows how to play the piano? Surely, it could only be on the basis of available evidence. Given our modification to Ryle's account of dispositions we can say that we would deduce what she knows from the 'picture' we have of her, construed in turn from what we have gleaned about her history, her years of training, her past performances, and so on. Again, insofar as we understand Ryle as saying that our judgements about what other people know are necessarily based on outward manifest evidence his case still stands.¹⁰

Here is another attempt to question the relation between knowledge and its manifestations, one that has been widely cited in the literature:

Suppose a famous dancer was to perform before an audience, an item from his repertoire to which he has himself given the following title:

A performance of Improvisation No. 15

To the astonishment of a member of his audience who just happens to be an expert on communications, the movements of the dancer turn out to resemble an accurate (movement perfect) semaphore version of Gray's 'Elegy', though the dancer is quite unaware of this fact.

(Carr, 1979, p. 407)

Yet again, understood as a question of how and what it is possible to know about what the dancer knows there is no mystery. On the evidence of the performance alone the communications expert might assume, wrongly, that the dancer knows semaphore and that the allusion to Gray's *Elegy* is intentional. Someone who is in a position to draw on a broader range of evidence would judge otherwise. By definition, judgements of significance always provide a better indication of what a person knows over any judgement of identity because they are able to draw on any or all available evidence. It is precisely by this means that we are able to discern fluke performances or performances not representative of a

person's true abilities, instances of cheating, and so on. Again, none of this does anything to undermine Ryle's point about our lack of access to other minds and his claim, in effect, that our judgements about what other people know are necessarily based on outward manifest evidence. And it is worth noting how so many of the examples used to illustrate this issue – the often bizarre references to such things as stunt skiers, pianists (or chefs) with no arms, guitar teachers unable to play well, pilots who can fly but don't know how to fly, and so on – can be seen to conflate the two kinds of question at issue here, that is, questions about the nature of knowledge, and questions about how and what it is possible to know about what a particular person knows.

So where does this leave us as regards knowing how and knowing that? Certainly, the terms have application in circumstances where it is sufficient to denote the 'overt acts and utterances' that are the outward manifestations of knowing. Associating our third person knowledge attributions with specific acts or utterances often serves to convey more clearly what it is a person knows. In informal or non-critical situations these reports of judgements of identity – for that is essentially what they are – will often be sufficient. However, such attributions are always susceptible to being trumped by reports based on judgements of significance, particularly in circumstances where high stakes are attached to those judgements, and judgements of this kind are *not* amenable to being couched in these terms. In such circumstances we typically abandon references to knowing how and knowing that in favour of terms such as 'knowledge', 'understanding', 'comprehension', 'grasp', and so on – notions which are characteristically un-bifurcated and clearly *not* associated with specific manifestations. Indeed, in such situations we may effectively *retract* attributions of knowledge implicit in our reports of a person's acts and utterances. The teacher of our earlier example would say of her pupil: 'He gave the correct answer but he didn't really know/understand/comprehend/grasp'. She would not say 'He knew that x but he didn't really know that x ' or 'He knew how to y but he didn't really know how to y '.¹¹ This is not a point about grammar or language use; it is rather that attributions of knowledge derived from judgements of significance always take precedence over attributions based on judgements of identity since by definition they draw on a wider range of evidence.¹²

Notwithstanding the above point, the notions of knowing how and knowing that can sometimes have a role in connection with judgements of significance. I might deduce from the 'picture' I have of a person that they would probably know how to dismantle a carburettor

or know that the capital city of Australia is Canberra, even though to I have no knowledge of them ever doing or saying either of these things. In other words, we can indicate something of what a person knows by indicating how their knowledge *might* be made manifest.

Importantly, however, I could never properly represent what I know of another person's mind in these terms because the 'picture' that enables me to make such knowledge attributions is not itself bifurcated in this way. It will be in large part indistinct and amorphous, an amalgam of half-remembered histories, diverse episodes and variously gleaned impressions. But it is none the less important for that, for ultimately it is by this means that we are able to make the very best estimation of what another person knows.

We might say, then, that the notions of knowing how and knowing that have entirely viable uses in the context of third person knowledge attributions. In this much Ryle was correct. Importantly, however, not all of our knowledge attributions can be couched in these terms. While judgements of identity are necessarily associated with specific behavioural manifestations this is not the case with judgements of significance. Indeed, whenever it is important to provide the fullest account of what a person knows we almost invariably abandon our use of such terms in favour of an account by which we purposely attempt to describe knowledge *rather than* its manifestations. Nowhere is this more evident than in circumstances in which judgements of identity (e.g. 'He was able to do x...') are countermanded by judgements of significance (e.g. '...but he didn't know what he was doing').¹³ There is thus an important sense in which the more serious we are about our attributions of knowledge the less resort we will have to the notions of knowing how and knowing that. This, I want to suggest, explains how the KH/KT distinction sits in relation to questions about the attribution of knowledge. But it is when we turn to the *other* kind of question, the sort of question concerned not with what a particular person knows but with how we should properly characterise the knowledge at issue that the more serious limitations of the KH/KT distinction become apparent.

A QUESTION OF KNOWLEDGE

While in some places, such as with his illustration of the clown/clumsy man, Ryle is concerned with the question of how and what it is possible for us to know about what another person knows – for example, how we are able to 'recognise that a performance is an exercise of a skill' (1949, p. 33) – in other places he is clearly concerned with questions of the second kind, that is, questions about the nature of knowledge. And it is in this vein that Ryle

challenges the intellectualist legend by which account, according to Ryle, it is assumed that a chef ‘must recite his recipes to himself before he can cook’, or the hero ‘lend his inner ear to some appropriate moral imperative’ (p. 29) before being able to save a drowning man, and so on. The question here, then, is about the kind of knowledge involved in acting in such and such a way.

The KH/KT distinction enters into Ryle’s attack on the intellectualist legend by way of his trying to show that the knowledge required for intelligent action does not consist in the prior consideration of propositions, rules, maxims, principles, etc. Now it would have been entirely possible for Ryle to make this case without invoking the KH/KT distinction at all; he might reasonably have offered an account of intelligent action in which the prior manipulation of propositions, rules, maxims, etc., could be seen to play little or no part. The apparent advantage for Ryle, in characterising such prior mental operations as ‘knowing that’ as opposed to ‘knowing how’ is that it lends *prima facie* plausibility to his case: if we accept that the kind of knowledge required for intelligent action is a knowing how as opposed to a knowing that, then it would seem to follow that anything characterised as knowing that is simply the wrong kind of knowledge. It is thus that Ryle effectively becomes committed to knowing how and knowing that as *epistemological* categories rather than merely the means by which we indicate how a person’s knowledge is made manifest. And of course a good many since Ryle have followed him down this very same path, accepting at face value this dual epistemic scheme of things.

Certainly some have *claimed* to repudiate the distinction. Stanley and Williamson purport to ‘contest the thesis that there is a fundamental distinction between knowledge-how and knowledge-that’; yet this claim is immediately thrown into question when they add: ‘Knowledge-how is simply a species of knowledge-that’ (2001, p. 411). For to suggest that A is a ‘species’ of B is not to deny the existence of A. It is not to say that when we refer to instances of A we are mistaken. Neither is it to say that we are mistaken in distinguishing between A and B. Rather, to claim that one kind of knowledge is a species of the other is to suggest that one stands in a certain relation to the other, or more precisely, that one has a certain kind of priority over the other. Similarly with those who claim to reduce the distinction in the other direction, such as with Hetherington’s contention that ‘To know that *p* is to know how to perform various actions’ (2006, p. 72). Hetherington is right when he

characterises the intellectualist position as affording ‘conceptual centrality’ (*ibid.*) to knowing that, but by the same token he merely assumes the same sort of ‘centrality’ for knowing how.

None of this conceptual gerrymandering does anything to dissolve the KH/KT distinction but it does testify to the fundamental instability of these supposed concepts in the context of questions about the nature of knowledge. This instability arises from the fact that each might variously be conceived in either conceptually ‘thin’ or ‘thick’ terms: either as little more than perfunctory act or utterance, or something that is epistemologically and cognitively far richer.¹⁴ Which way the priorities run is an entirely arbitrary matter, dependent only the prior conceptual commitments or predilections of the commentator. There is no means of adjudicating between thin/thick and thick/thin conceptions of knowing how and knowing that because the assumption upon which both are founded is mistaken. It is simply *assumed* that the distinction has some substantive epistemological basis. The mere fact that we distinguish between acts and utterances as a means of differentiating a person’s knowing one thing and knowing something else is taken to imply two kinds of knowledge. The contingent facts of language use are often recruited to bolster this assumption, yet it is clear that when *attributing* knowledge, we say ‘S knows how to x’ or ‘S knows that y’ we mean nothing other than what S knows – whatever that might be – is such as would enable S to act in such and such a way or produce such and such an utterance. We might equally determine that what S knows is such as could manifest itself both in certain acts *and* in certain utterances. But the essential point here is that our use of such locutions *in the context of knowledge attributions* neither entails nor necessitates two kinds of knowledge.

Those who have been keen to show that know-how can obtain in the absence of outward manifestations, or vice versa, have missed full the significance of this fact. Far from establishing the priority of one form of knowing over the other it should more properly be regarded as demonstrating the important sense in which knowledge cannot be conceived in terms of its manifestations and thus in terms of either knowing how or knowing that.

Consider the following by no means unusual scenario:

A factory production line is in full swing when suddenly the machines grind to a halt. Alarm bells ring and warning lights flash; a maintenance technician arrives and makes his way to one of a hundred electrical control panels each interconnected perhaps with several miles of cabling. He opens the control panel, takes a screwdriver from his pocket and makes a small adjustment to just one of several hundred components. Closing the control panel he presses some buttons and the production line bursts into life. The question is, how is it possible to account for what the technician knows? His performance did not require the conscious manipulation of propositions or facts – and neither did it require any particular

physical dexterity.

(Lum, 2009, p. 56)

What we *can* say is that the knowledge requirements here are considerable and such as would typically require several years' formal classroom-based provision and extensive off-the-job training, not to mention a good deal of post-qualification experience. If we wished merely to make an *attribution* of knowledge we might say that the technician 'knows how to reset the overload relay' – the 'knows how' locution serving to indicate that he has the wherewithal to effect the *doing* of something. The difficulty with this locution – centred on the simple physical act of pressing a button – is that it fails quite radically to identify the knowledge at issue. To see the extent of this failure, suppose that a machine operator who had previously observed the technician at work took it upon himself to ape the technician's actions. Of course he would have no way of knowing whether the action was appropriate in this particular case, still less what its consequences might be. With no understanding of what he was doing he would have no idea whether further investigation was needed; he would be unaware of the risks posed to equipment, his safety and that of others. What the machine operator would lack is not merely some part of what the technician knows but virtually *all* of what he knows. Yet it would still be true to say, in the judgement of identity sense, that the machine operator 'knows how to reset the overload relay'.

Now it might be thought that we could give a better account of what the technician knows by adding further attributions couched in these terms. Perhaps we could say that unlike the machine operator the technician also 'knows *how* to distinguish different kinds of faults', 'knows *that* it is dangerous to..., ' and so on. This much is true and certainly on this basis we could easily determine that the technician knows things that the operator does not know. Yet this would be to miss the point, for the question here is not how we might *assess* what the technician knows but, rather, *what* it is the technician knows that equips him to do the task. And the difficulty here is that there is simply no evidence that any such 'extra' acts and utterances were involved in his performance, either overtly or even – and Ryle was surely correct on this point – as private, silent performances in the technician's mind.

The crucial thing here is to recognise how our acts and utterances are *derivative* rather than constitutive of what we know. Take, for instance, someone who drives for a living and has come to have a good knowledge of central London. They are able to find their way around, recognise landmarks, plot in their mind how best to get from A to B, etc. Suppose it to be

true of this person that they know *how* to get from Camden to Charing Cross, that they know *that* The National Gallery is in Trafalgar Square, and so on. If we wanted to assess what the driver knows we could imagine testing them on a selection of just these kinds of things. But this is not to say that what the driver knows is the sum total of some such list of acts and utterances. Importantly, such a list could never be complete because it would always be possible to elicit further acts or utterances. If asked, he could perhaps tell us whether the stonework of Westminster Abbey is lighter or darker in colour than St Pauls, even though the thought had never before crossed his mind. And his knowledge would enable him to undertake any number of journeys he has never made before. Indeed, what the driver knows could give rise to a virtually infinite number of acts and utterances. There is thus an important sense in which any specification couched in terms of knowing how and/or knowing that will necessarily underdetermine the knowledge that is substantively at issue.

We get a sense of the extent of this shortfall by considering the knowledge we bring to bear in our use of language. Take, for example, the sentence ‘She gave him her key and he opened the door’. As John Searle (1995) says, in linguistic pragmatics it is a moot point whether it is actually said or just implied that the door was opened with the key, that the door was opened *after* she gave him the key, and so on. Either way, there is broad agreement ‘that there is a certain underdetermination of what is said by the literal meaning of this sentence’ (p. 131). But in fact this, on Searle’s view, is to understate the matter:

I wish to say there is a *radical* underdetermination of what is said by the literal meaning of the sentence. There is nothing in the literal meaning of the sentence ... to block the interpretation, He opened the door with her key by bashing the door down with the key ; the key weighed two hundred pounds and was in the shape of an axe. Or, He swallowed both the door and the key and he inserted the key in the lock by the peristaltic contraction of his gut.

(Searle, 1995, p. 131)

According to Searle, our ability to understand even the simplest linguistic construction is dependent upon our having an extensive ‘Background’ understanding of how the world works, an understanding that is presumed by rather than contained in the constructions themselves. This is significant for our purposes because it would seem to suggest that there is a certain kind of epistemic deficit with any designation of the form ‘S knows that *p*’ in that it omits the crucial matter of what S needs to know in order to understand and make sense of *p*. Similarly, to say that ‘S knows how to *x*’ is merely to denote one manifest outcome of what S knows while leaving unanswered the question of *what* it is S knows.

Elsewhere (Lum, 2009) I have drawn attention to parallels between Searle's notion of Background and approaches in phenomenology which emphasise the intentionality or directedness of knowledgeable states.¹⁵ Probably the fullest expression of this approach is to be found in Martin Heidegger's (1962) famous 'analytic of *Dasein*', the 'there-being' of human existence which, I have suggested, allows a far more coherent account of knowledge and expertise than is possible within the Rylean scheme of things. There is not space here to rehearse such an account in detail but suffice it to say that on this view, to learn or to become expert is first and foremost about the learner having disclosed to them a 'world' of interconnected meanings and involvements; we might, for example, think of the 'worlds' of mathematics, literature, fine art, engineering, motherhood, medicine, playing chess, and so on. On this view, what the technician and the driver have in common is that each is able to make sense of the particular 'world' in which they operate. Each has an understanding of the priorities and purposes of that world, is able to cope and find his way around in it, recognise things *as* certain things and understand the connections between them. Each will have appropriate expectations, an understanding of possible scenarios, likely outcomes, and so on. In the absence of such knowledge, acts and utterances are no more than mechanical behaviours, such as might be learnt by rote or merely parroted.

We should have no objection to referring to this knowledgeable coping as 'know how' were it not for the fact that this will be taken by some to imply one particular *kind* of knowledge: i.e. a knowing how as opposed to a knowing that. Importantly, this same knowledge is fundamental *both* to purposeful action and the meaningful use of propositions. So neither could we feasibly conceive of this knowledge as a 'knowing that'. Indeed, given that the epistemic content of any 'knowing that' construction is in large part constituted of this knowledge, to characterise it as 'knowledge that' would simply invoke a regress whereby in order to 'know that *x*' we would need to 'know that *y*', and in order to 'know that *y*'... etc. Certainly, part of what is known will involve knowing *how* things should be done but neither this, nor the fact that we might couch this in terms of knowing *that* there is a way to do something, does anything to support the dual knowledge thesis or the claim that one ostensible form of knowledge is a species of the other. That the knowledge at issue cannot readily be accommodated into an account centred on acts and utterances might, for some, suggest a *third* kind of knowledge, something akin to acquaintance.¹⁶ But on the view presented here, far from being additional to the knowledge which enables us to do and say things, knowledge gained by acquaintance should properly be regarded as part of the doxastic

ground, i.e. the complex of beliefs, which enables us to produce meaningful acts and utterances. In this sense, knowledge gained from acquaintance underpins *both* my ‘knowing that’ this apple is red, and my ‘knowing how’ to tighten a screw the right amount, which again militates against the feasibility of the assumption that that these two terms represent two distinct forms of knowledge.

Ryle was right to dismiss as mistaken the idea that to perform intelligently is ‘to do a bit of theory and then to do a bit of practice’ (1949, p. 29). The terms theory and practice, like knowing how and knowing that, can be used to indicate how a person’s knowledge is made manifest, such as when we say of someone that they ‘know the theory’ or ‘know the practice’ of something, again – at the risk of labouring the point – by way of indicating the difference between knowing one thing and knowing another. But the terms ‘theoretical’ and ‘practical’ have another connotation in education, for they are also used to signify the *manner* of learning, such as when we want to differentiate learning from a text from learning by means of a practical exercise.¹⁷ Again, the distinction is not clear-cut, for just as not all manifestations of knowledge can be categorised as being either act or utterance, similarly with the separation of provision into the theoretical and the practical. The best ‘theoretical’ provision may sometimes be that which has the most acute practical relevance, the best ‘practical’ provision that which most enhances the learner’s understanding of things. Certainly we can expect different things to be learnt from different modes of provision. The mistake is to assume that this can be taken as signifying two different kinds of knowledge. Knowledge, I want to suggest, can no more be categorised by its sources than by its manifestations.¹⁸

CONCLUSION

The upshot of this is that although the KH/KT distinction has some limited viability in the context of knowledge attribution it is simply not viable as an epistemological distinction. This is by no means merely of theoretical import for the idea that there are two fundamentally different kinds of knowledge can be seen to be of very real consequence for education. Nowhere is this more evident than in vocational and professional education where the effects this dichotomous way of thinking can be seen to have a number of profoundly negative effects. In the context of assessment, the association of third person knowledge attributions with explicit acts and utterances has an irresistible bureaucratic and managerial appeal. It allows educational achievements to be specified in exact and explicit terms: the learner will

‘know how to ...’ or ‘know that ...’ and couched in these terms the assessment process appears deceptively simple, requiring only simple judgements of identity. This is, of course, precisely the thinking that lies behind the current vogue for ‘outcomes’, ‘competencies’ and ‘skills’. The difficulty, as we have seen, is that such judgements are simply inadequate in high stakes circumstances when it is vital to determine what learners actually know. First, because arrangements of this kind cannot provide the best estimation of knowledge, and second, because such arrangements can have the effect of requiring assessors to record results in the face of contrary indications and against their better judgement.¹⁹

But it is in connection with matters of provision and curriculum design that this dichotomous conception of knowledge stands to have the most detrimental effects. It hardly needs to be said that both theoretical and practical modes of provision have an indispensable part to play in the development of occupational expertise. Difficulties arise when it is assumed that the purpose of each is to provide a different kind of knowledge. First, because this is to lose sight of the important sense in which *both* kinds of provision should contribute to the knowledge that is substantively at issue. To lose sight of this is to risk theoretical provision drifting into irrelevance, theory for theory’s sake, and practical provision being reduced to instilling rote behaviours. Second, it becomes a moot point which of the two ostensible forms of knowledge should have priority. While some will insist that the facility to act is all important, to know *how* to ‘do the job’, others will be adamant that it is ‘theory’, knowing *that*, that is the wellspring of expertise. In the UK we need look no further than the training of nurses and teachers to witness the ill effects of this dichotomous conception of knowledge; while nurse education has moved from the hospital ward to the university, teacher training is purposefully being shifted in the opposite direction, away from the university to the workplace. Such confusions at the level of national policy betray deep-seated attachments to the dual knowledge thesis and erroneous assumptions about the relative importance of one ostensible form of knowledge over the other. The truth is, of course, that properly conceived both modes of provision should have a vital role in developing the capabilities at issue. Indeed, perhaps the most serious danger posed by dual knowledge thesis is that it could encourage the thought that it is possible to dispense with one form of knowledge altogether. There are no doubt many who would concur with David Carr when he concludes his analysis of ‘knowledge in practice’ with the following observation:

A profound mistake is made in supposing that the best way to turn an apprentice into a master plumber is to instruct him in the complete theory of hydraulics, for the practical knowledge that the apprentice requires comes with initiation into and mastery of practical rather than theoretical discourse.

(Carr, 1981, p. 61)

The difficulty with this is that it simply does not square with reality. It risks grossly misrepresenting the nature and extent both of the knowledge required and of the enterprise needed to provide that knowledge. It encourages the thought that the apprentice has no need of formal ‘theoretical’ provision, no need of systematised knowledge, and that the emphasis should properly be on doing rather than thinking.

The longstanding assumption that there are two kinds of knowledge, a ‘knowing how’ and a ‘knowing that’ is not only epistemologically incoherent but potentially of profound detriment to education. Although there is little evidence that education policy makers of late have ever taken their lead from philosophy, it would be unrealistic to expect education to dispense with ill-conceived and unfeasible ways of thinking about knowledge if those same ways of thinking continue to pervade the philosophic literature.

Correspondence: Gerard Lum, School of Education, Communication & Society, King’s College London, Stamford Street, London SE1 9NH, UK.
Email: gerard.lum@kcl.ac.uk

NOTES

1. Ryle himself acknowledges that the kind of understanding at issue ‘does not consist in inferring, or guessing, the alleged inner-life precursors of overt actions’ (1949, p. 54). Indeed, he emphasises repeatedly that the process is not one of inference.
2. See, for example, M.L.J. Abercrombie’s (1989) classic account of empirical work in this area, demonstrating the important role of judgement in perception.
3. It would be a mistake to characterise the difference between these two forms of judgement in terms of objectivity/subjectivity; in education there is often resort to this way of thinking. Rather, the difference between the two approaches consists in the stance taken towards the evidence, either confirming the presence of predetermined evidence, or alternatively being open to any and all relevant evidence. To characterise the latter as ‘subjective’ is to miss the point that judgments based on the fullest range of evidence *always* provide the best indication of what a person knows (see Lum, 2012).
4. On this point see Rorty (1980) and Lum (2009). Ryle is far from consistent on this score, as when he declares: ‘Of course it is part of my general thesis that the supposed occult processes are themselves mythical; there exists nothing to be the object of the postulated diagnoses’ (Ryle, 1949, p. 54), thus perhaps leaning more towards a potential interpretative role for the observer.
5. This is presumably what Bengson and Moffett (2011a) have in mind when they note: ‘What is distinctive of anti-intellectualism is its commitment to the thesis that knowing how requires the *corresponding* ability or disposition’ (original emphasis; p. 167).
6. Although widely used in this connection in the literature, the term ‘ability’ is unhelpfully prone to ambiguity. On one reading, ‘ability’ might reasonably be interpreted as indicating that which equips a person to perform in a particular way, the wherewithal to act, so to speak. In contrast, what is at issue here are the ‘overt acts’ that are the outward manifestations of knowing.
7. Bengson and Moffett (2007), for example, describe it as ‘rather puzzling’ (p. 32).
8. The failure to acknowledge this simple point is the cause of no end of puzzlement in the literature. Bengson and Moffett (2011a), for example, contrast Pat (who cannot ski but can teach ski stunts) with Albert (who

similarly cannot ski but ‘knows the theory’ of skiing). On Bengson and Moffett’s analysis ‘...Pat and Albert both know how one does the stunts; neither is able to do the stunts. *But plainly a significant difference remains: only Pat knows how to do the stunts*’ (original emphasis, p. 169). It is perhaps somewhat disingenuous to suggest first that both Pat and Albert know the same thing in virtue of knowing ‘how one does the stunts’ only then to declare that there is a ‘*significant difference*’ between what each knows. This manoeuvring is justified by reference to supposed differences between such locutions as ‘knowing how *one* ϕ -s’ and ‘knowing how *to* ϕ ’. A more coherent analysis would simply acknowledge that what each knows is different. Pat knows how to teach others to do stunts and Albert knows how to make sense of, perhaps recite, certain theoretical information about skiing – and each is different again from someone who could actually perform the stunts. Certainly if Pat knows some of the theory that Albert knows then to that extent they could be said to know the same thing. But beyond this what each knows is obviously very different and it is misleading to say that they ‘both know how one does the stunts’.

Similarly with Stanley and Williamson’s attempt to construct Gettier-style counterexamples for instances of knowing how:

Bob wants to learn how to fly in a flight simulator. He is instructed by Henry. Unknown to Bob, Henry is a malicious imposter who has inserted a randomising device in the simulator’s controls and intends to give all kinds of incorrect advice. Fortunately, by sheer chance the randomising device causes exactly the same results in the simulator as would have occurred without it, and by incompetence Henry gives exactly the same advice as a proper instructor would have done. Bob passes the course with flying colors. He has still not flown a real plane. Bob has a justified true belief about how to fly. But there is a good sense that he does not *know* how to fly...

(S and W, 2001, p. 435)

Yet it is clearly *not* the case that the presence of the randomising device undermines Bob’s claim to knowledge. Presuming that he performed just as any knowledgeable pilot would have done then the presence of the device is neither here nor there. A person’s knowledge of how to drive a car is not thrown into question by a malicious but failed attempt to tamper with the car’s controls. And if that person learned to drive in a car that *had* been altered then they would know how to drive *that* car. Putting aside the apparent suggestion that knowing how to fly somehow equates to following advice, again we can see this as a case of failing to recognise the difference between knowing one thing and knowing another. There is clearly a difference between knowing how to: (i) follow advice (whether correct or incorrect), (ii) distinguish correct from incorrect advice, (iii) fly correctly (in simulation) given certain correct advice, (iv) fly correctly without need of advice ... etc.

9. Confusions arising from this kind of ambiguity abound in the literature. Take, for example, the following:

A guitar teacher might know how to play well, but lack the ability to play well. Perhaps he’s uncoordinated or never practices. Further, his ability to play might improve with practice, even as his know-how remains constant. So, his know-how doesn’t reduce to the corresponding ability.

(Alter, 2001, pp. 232-3)

In what possible sense could the guitar teacher be said to ‘know how to play well’ if he cannot play well? Are we to take it that by ‘know how’ is meant the knowledge needed to teach others to play well, or is it perhaps knowledge of what *he* should do in order to play well? But then the claim that *this* know-how ‘does not reduce to the corresponding ability’ appears incoherent, if by ‘ability’ is meant his ability to play the guitar rather than, say, teach it.

10. We can take it that without this corrective Ryle would have difficulties with a case such as this.

11. Compare this with Bengson and Moffett’s (2007) claim that ‘there is no reading of the following sentence on which it is not contradictory: Irina knows how to do a quintuple salchow, but she doesn’t know how to do a quintuple salchow’ (p. 39). In point of fact there *are* readings which would render this sentence non-contradictory. The difficulty with a construction such as this is not that it is necessarily contradictory but that it fails to convey what it purports to convey: the sense in which Irina does know and the sense in which she does not know.

12. The point here is that judgements of significance include the evidence upon which judgements of identity are based *in addition to* whatever other evidence is deemed significant.

13. This in part explains what Bengson and Moffett (2007) see as the ‘peculiar behaviour of know-how attributions’ (p. 40), why they sometimes entail ability attributions and sometimes do not. Bengson and Moffett’s proposed solution to this ‘puzzle’, centres on ‘concept possession’ (2007, p. 55) on the part of the knower. But *pace* Bengson and Moffett, we can see that the matter hinges not on the presence or otherwise of any particular epistemic characteristic but, rather, on the kind of judgement upon which attributions of knowledge are based. Moreover, as I have explained elsewhere, contrary to longstanding traditions in analytic philosophy there are profound difficulties with resorting to the notion of ‘concepts’ understood as discrete constituents of knowledge (see Lum, 2015).

14. Accordingly, while some, such as Ryle, adopt a conceptually thin notion of knowing that, others opt to reverse this priority, characterising knowing how as mere behaviour and portraying knowing that in conceptually thick terms (for an extended account this see Lum (2009) ch. 3)

15. In Lum (2009) I develop at length an account of occupational expertise which draws on diverse theoretical perspectives in analytic philosophy, phenomenology and psychology of perception.
16. The notion of knowledge by acquaintance has a lengthy provenance from Plato through William James to Bertrand Russell who distinguishes between knowledge by acquaintance and knowledge by description according to their source. For a recent account of the educational implications of knowledge by acquaintance understood as a third form of knowledge see Winch (2013).
17. I have suggested elsewhere that we might usefully refer to the modes of learning and the manifestations of knowing respectively as the antecedent and consequent conditions of learning (see Lum, 2007).
18. It is patently the case that knowledge gained from a textual or theoretical source can manifest in the ability to do something practical, and equally, knowledge gained by practical or experiential means can manifest in a person's being able to say certain things, answer questions, etc.
19. The point here is that high stakes assessment should properly employ judgements of significance which are often not amenable to being couched in terms of acts and utterances, i.e. knowing how and knowing that.

REFERENCES

- Abercrombie, M. L. J. (1989) *The Anatomy of Judgement* (London: Free Association Books).
- Alter, T. (2001) Know-how, Ability, and the Ability Hypothesis, *Theoria*, 67, 229-239.
- Bengson, J. and Moffett, M.A. (2007) Know-how and concept possession. *Philosophical Studies*, 136, pp. 31-57.
- Bengson, J. and Moffett, M.A. (2011a) Non-propositional Intellectualism, in J. Bengson and M.A. Moffett (eds), (2011b) *op cit*, pp. 161-195.
- Bengson, J. and Moffett, M.A. (eds) (2011b) *Knowing How: Essays in Knowledge, Mind and Action* (Oxford: Oxford University Press).
- Carr, D. (1979) The Logic of Knowing How and Ability, *Mind*, 88, pp. 394-409.
- Carr, D. (1981) Knowledge in Practice, *American Philosophical Quarterly*, 18, 1, pp. 53-61.
- Davidson, D. (2001). Knowing one's own mind. In *Subjective, intersubjective, objective* (pp. 15-38) (Oxford: Oxford University Press).
- Heidegger, M. (1962) *Being and Time*. (Oxford: Blackwell).
- Hetherington, S. (2006) Knowledge-that is Knowledge-how, in S. Hetherington (ed) *Epistemology Futures* (Oxford: Oxford University Press).
- Lum, G. (2007) 'The myth of the golden mean', in J. Drummond and P. Standish (eds.) *The Philosophy of Nurse Education* (London: Palgrave Macmillan).
- Lum, G. (2009) *Vocational and Professional Capability: An Epistemological and Ontological Study of Occupational Expertise* (London: Continuum).
- Lum, G. (2012) Two Concepts of Assessment, *Journal of Philosophy of Education*, 46, (4) 589-602.
- Lum, G. (2015) Afterword: Can the two positions be reconciled, in A. Davis, C. Winch and G. Lum (ed) *Educational Assessment on Trial* (London: Bloomsbury).
- Rorty, R. (1980) *Philosophy and the Mirror of Nature* (Oxford: Blackwell).

Ryle, G. (1949) *The Concept of Mind* (London: Hutchinson).

Searle, J. R. (1995) *The Construction of Social Reality*. (London: Penguin Press).

Snowdon, P. (2004) Knowing How and Knowing That: a Distinction Reconsidered, *Proceedings of the Aristotelian Society*, 104. 1, pp. 1-29.

Stanley, J. and Williamson, T. (2001) Knowing How, *Journal of Philosophy*, 98.8, pp. 411- 444.

Winch, C. (2013) Curriculum Design and Epistemic Ascent, *Journal of Philosophy of Education*, 47, 1, pp. 128-146.