

Expanding the Client's Perspective

Yuri Cath

La Trobe University

The Philosophical Quarterly

[This is a preprint of an article whose final and definitive form is published in *The Philosophical Quarterly* 2023, online first <<https://doi.org/10.1093/pq/pqad032>>]

Abstract: Hawley introduced the idea of *the client's perspective* on knowledge, which she used to illuminate knowing-how and cases of epistemic injustice involving knowing-how. In this paper I explore how Hawley's idea might be used to illuminate not only knowing-how, but other forms of knowledge which, like knowing-how, are often claimed to be distinct from mere knowing-that, focusing on the case studies of moral understanding and 'what it is like' knowledge.

Keywords: understanding, experiential knowledge, the client's perspective

Introduction

If we didn't have a concept of knowledge, why would we invent it? What basic human needs does this concept serve? Craig (1990) famously wanted to reorientate epistemology towards such questions, in the hope that we might gain a better understanding of the nature of our concept of knowledge by reflecting, first, on what its function is and then, after that, thinking about what our concept of knowledge would need to be like in order for it to fulfil that function. And Craig's influential hypothesis about the function of our concept of knowledge is that it serves our shared interests as *inquirers* in identifying reliable informants.

One challenge for Craig's view is how well it extends to our thought and talk about knowing-how. For when we ascribe knowing-how our focus seems to be on whether the subject in question is a reliable performer, rather than a reliable informant. With such issues in mind, Hawley (2011) introduced the idea that our concept of knowing-how often serves our interests not as inquirers but as *clients* in finding people who can reliably perform certain actions for or with us. And Hawley explored how this idea might shed new light on knowing-how, and cases of epistemic injustice involving knowing-how.

My aim in this paper is to explore how Hawley's idea might usefully illuminate not only our thought and talk about knowing-how, but also much larger patterns in how we think and talk about different forms of knowledge and understanding. In §1 I review Hawley's discussion of the client's perspective and how it contrasts with the inquirer's perspective, and I will show how these ideas point us towards a distinction between two kinds of epistemic states that I call 'upstream knowledge' (related to the inquirer's perspective) versus 'downstream knowledge' (related to the client's perspective). In §2 I will discuss Hawley's brief but useful remarks on how the client's perspective can apply to forms of knowing-wh other than knowing-how, and I provide a kind of template for how the concept of downstream knowledge could be extended to these and other cases. I turn then in §§3-4 to my main task which is to argue that the client's perspective, and the idea of downstream knowledge, can illuminate key features of two other kinds of epistemic states that, like knowing-how, are often claimed to be distinct from mere knowing-that, namely, understanding-why and 'what it is like'-knowledge. The features in question being that, in certain circumstances, relevant forms of knowing-that appear to be neither sufficient nor necessary for being in these epistemic states. In §5 I close by briefly considering some linguistic and conceptual concerns one might have about these ideas.

1. Clients and Streams

We all have perceptual and reasoning faculties for gaining information about our world. But the amount of information we can each gain from our own 'on-board' faculties is limited. As inquirers then we have an interest in gaining access to the information that other people have gained through their faculties. If, say, I am considering whether to go to into a forest I have never ventured into before, and you have been there recently, I may have an interest in acquiring any information you have about whether there are any dangerous animals there, or whether there is any food to forage. And often our only way of accessing such information from other people will be through their testimony. But people can lie and have false beliefs, so our interest as inquirers in gaining information from other people leads us to an interest in identifying reliable informants. And Craig (1990) suggests that the function of our concept of knowledge is to serve our shared interests as inquirers in identifying reliable informants.

1.1 The Client's Perspective

Craig (1990: 156) tries to account for the function of our concept of knowledge-how in a similar way by introducing the notion of the *apprentice*—which is meant to closely parallel that of the inquirer—who seeks a person who can tell or show them how to Φ because they desire to know

how to Φ themselves. But, as Hawley points out, sometimes we seek people who know how to Φ not because we wish to gain this knowledge for ourselves (as we may possess it already or just not be interested in acquiring it), but rather because we seek someone who can reliably perform actions of Φ -ing for or with us:

There is, however, a further kind of motive for seeking someone who knows how, a motive that may be very central to our thinking about knowledge-how. When I seek a plumber, hairdresser, or architect, usually this is because I need the drains fixed, my hair cut, or a building designed. I need have no interest in learning how to do these things myself, nor in finding someone who can either teach or assess others. Perhaps I know how to do such things already but am too busy or too lazy to get them done myself (and I can't reach to cut my own hair). I call this 'the client's situation,' in contrast with the inquirer's and the apprentice's situations. (Hawley 2011: 287)

If we assume that this client's situation or perspective is "very central to our thinking about knowledge-how" what views of the function and nature of our concept of knowledge-how emerge from that assumption?

Minimally, the resulting concept will be one on which there are tight connections of some kind between knowledge-how and ability, given the client's interest in identifying reliable performers. One might even suggest that the client's perspective supports a simple anti-intellectualist view on which knowing how to Φ is (i) solely a matter of having an ability to Φ , and (ii) not a matter at all of being in any kind of propositional attitude state. But, as Hawley discusses, challenges have been made both to the claim that knowing how to Φ entails having an ability to Φ (Stanley and Williamson 2001) and to the claim that having an ability to Φ entails knowing how to Φ (Snowdon 2004).

In the literature, however, there have now been numerous arguments that provide a strong case that related entailment theses may still hold, and in both directions, once one restricts them not to the mere ability to Φ but, following Hawley (2003), to something like the ability to Φ *intentionally in normal circumstances*.¹ And, as Hawley (2011) points out, the client would be interested in identifying someone who has not just a 'bare ability' but someone "who can control her ability, exercising it at will, repeatedly, and in line with the client's wishes" and

¹ For discussions that support the claim that having an ability to Φ intentionally entails knowing how to Φ see Cath (2011, 2015a), Constantin (2017) Hawley (2003), and Löwenstein (2017). For discussions that support the claim that knowing how to Φ entails having an ability to Φ intentionally see Cath (2015a; fn.14), Glick (2012), Hawley (2003), Löwenstein (2017); and Pavese (2015). For a defence of both entailment theses see Pavese (2021).

that knowledge-how “may well be necessary for this kind of intentional, responsive, controlled ability” (*ibid*: 288).

Should we adopt a qualified anti-intellectualist view then on which knowing how to Φ is (i) solely a matter of possessing an ability to Φ *intentionally*, and (ii) not a matter of being in any kind of propositional attitude state? That is an important option but, I think, ultimately, an unattractive one even if we grant that knowing how to Φ entails having an ability to Φ intentionally and vice versa. This is because there are good arguments that having an ability to Φ intentionally itself requires one to have a true belief about a way or means of Φ -ing (see e.g., Pavese 2021), in which case even if we accept (i) this pushes us towards *rejecting* rather than accepting (ii).

For these reasons I think the concept of knowing-how that is best supported by Hawley’s work on the client’s perspective is one on which knowing how to Φ is a matter of being in a true belief state with the right kind of content (likely a proposition of the form ‘*w* is way for me to Φ ’), but this true belief state must also be possessed in such a way that being in it explains, and perhaps even entails, one’s ability to Φ intentionally.

1.2 Downstream Knowledge

With the above ideas in hand, we are now in a position to introduce the distinction between *upstream* versus *downstream* knowledge. The common core of these two forms of knowledge is true belief but then they differ with respect to what kinds of further properties are needed to ‘upgrade’ a mere true belief into a state of knowledge, with these differences relating to the differing interests of Craig’s inquirer versus Hawley’s client.²

Craig’s inquirer will be focused on what Hawley calls ‘upstream indicators’ of true belief, that is, properties of another person which the inquirer can detect, and which indicate that the way in which that person came to have their belief (hence, why they are ‘upstream’ properties), makes it likely to be a true belief. These could be properties like someone’s having been in the right location to come to have a true belief about whether *p* or their having a reliable method for forming beliefs about whether *p*.

Craig thought that by a process of ‘objectivisation’ (for discussion see Hannon 2019: 39-41) our collective interests in pooling information would’ve led us to our actual concept of

² I am relying here on the orthodox, but by no means uncontested, assumption that knowledge is a species of true belief. I suspect the upstream/downstream distinction could be redrawn in a knowledge-first framework that rejects this assumption, but I will not explore that issue here.

knowledge which is defined in terms of closely related but more objective upstream properties in the sense that they are not defined in terms of what is detectable by any one individual inquirer. And Craig ends up endorsing a reliabilist view on which the concept of knowledge is analysed in terms of true beliefs that are the product of reliable belief forming processes.

Note that there is a shift here not only from more subjective to more objective upstream properties, but also a shift from a notion of upstream *indicators* to a notion of upstream *conditions*. The former notion is epistemic as an upstream indicator is a property that gives the inquirer reason to judge that someone has a true belief, whereas the latter notion is metaphysical or conceptual as an upstream condition is a property a true belief needs for it to constitute knowledge.

Importantly, even if we do not accept Craig's own analysis of our concept of knowledge, we can note that almost all traditional analyses of knowledge rely crucially on different kinds of upstream, or 'backwards looking', conditions. So, for example, one's true belief being based on good evidence, or one's true belief not being subject to Gettier-style luck, or one's true belief being the product of an intellectual virtue, are all upstream conditions because they are all concerned with the etiology of one's belief in the sense of its causal or rational origins. As Hookway (2006: 105) notes, the "features that have been taken to be characteristic of knowledge have been backward-looking: they have concerned the history of the candidate belief or the kind of justification that the believer (already) possesses for it". To possess upstream knowledge then is to possess a true belief that meets certain upstream conditions, and when philosophers talk about 'knowledge' this is usually what they have in mind.

As opposed to Craig's inquirer, however, Hawley's client will be focused on what she calls 'downstream indicators' of true belief, which in the case of knowing how to Φ will typically be actions of successfully Φ -ing which are detectable by the client, and which the client will take to be indicators of the other person's having a true belief about how to Φ and, consequently, being a reliable performer of actions of Φ -ing. For example, if I see Stephanie opening the safe then I will likely conclude that she has a true belief about how to open the safe and that, consequently, she can reliably open the safe when she wants to.

Hawley's discussion is primarily focused on just the idea of downstream indicators of true beliefs about how to Φ , as opposed to downstream conditions for a true belief about how to Φ constituting knowledge how to Φ . But if we wanted to make the parallel shift that we identified in Craig's work, then the idea would be that our concept of knowledge-how itself is analysable

in terms of downstream properties related to those that the client looks for (but where these properties are more objective in the sense that they are not defined in terms of what one individual client can detect). However, we could not plausibly claim that someone would have to have, in fact, performed a given action in order to possess downstream knowledge as knowledge can obviously precede, and exist in the absence of, any performances guided by that knowledge. Relatedly, Hawley points out that “downstream indicators [i.e., successful actions] are more naturally thought of as consequences of knowledge, not constituents of it” (2011: 293).

The solution, of course, is to state the relevant downstream conditions in relation to the abilities themselves, rather than specific manifestations of those abilities. More precisely, the conditions will be characterised in the way outlined earlier, that is, in terms of a true belief state that explains, and perhaps even entails, a subject’s ability to intentionally perform certain relevant actions. So, Stephanie’s true belief that the way to open the safe is to enter the pin 4378 and then turn the handle to the right constitutes a state of downstream knowledge just in case her possession of that true belief explains her ability to perform possible future actions of opening the safe, even if she never actually exercises that ability.

Different accounts can then be given of what exact downstream properties true belief states will need to have in order for them to explain the possession of such abilities, including accounts which appeal to practical modes of presentation (Pavese 2015, Stanley 2011), and accounts which appeal to practical ways of being in a propositional attitude state (Waights Hickman 2019, Cath 2020), or accounts which claim that the propositional content alone entails the relevant ability (Stanley 2011, Cath 2017). I will not try to adjudicate among such accounts here, but I will assume that some such account (or combinations therefore) is correct.

To see more clearly how this distinction between upstream versus downstream knowledge works, take any true belief of the form ‘ w is a way for me to Φ ’. One has *upstream knowledge*, of some way w , that w is a way for oneself to Φ just in case (i) one has a true belief that w is a way for oneself to Φ , and (ii) this true belief has certain etiological properties. On the other hand, one has *downstream knowledge*, of some way w , that w is a way for oneself to Φ just in case (i) one has a true belief that w is a way for oneself to Φ , and (ii) this true belief is possessed in such a way that it explains one’s possession of the ability to Φ intentionally in normal circumstances.

In some cases, a subject's true belief that w is a way for them to Φ will meet both the relevant upstream and downstream properties and, hence, will constitute both upstream and downstream knowledge. But, in other cases, the subject's true belief will only constitute downstream knowledge and not upstream knowledge, or vice versa. And, in other work, I have appealed to this point—together with the hypotheses that knowing-that ascriptions are paradigmatically used to ascribe states of upstream knowledge, whereas knowing-how ascriptions are paradigmatically used to ascribe states of downstream knowledge—to help explain key intuitive judgments that have played central roles in debates about the nature of knowing-how (Cath 2015a, and forthcoming, and see Levy 2014 for related discussion). Specifically, I argue that these ideas can help us to explain both cases where knowing-that appears to be insufficient for knowing-how (these are cases where only upstream but not downstream knowledge is present), and cases where it seems not to be necessary for knowing-how (these are cases where downstream but not upstream knowledge is present). I won't walk through those arguments here, but what I want to show now is how the client's perspective, and the notion of downstream knowledge, can be applied in parallel ways to other forms of knowledge which, like knowing-how, are often claimed to be distinct from knowing-that.

2. Initial Expansions

As well as knowing-how, Hawley (2011: 288-289) briefly discusses how the client's perspective could apply to other forms of *knowing-wh*,³ in particular, knowing-whether, knowing-when, and knowing-where. In relation to knowing-whether, Hawley points out that knowing-how often involves knowing-whether, for example, knowing how to defuse a bomb involves knowing *whether* to cut the red wire or the blue. In which case, if the client's perspective is central to our thought and talk about knowing-how we might expect to also be important to how we think and talk about knowing-whether. Relatedly, Hawley notes that when we seek someone who knows *where* to buy cheap but reliable cars, or *when* to invest in the stock market, our interest may not be in acquiring such knowledge for ourselves, but just in finding an agent who can perform the associated actions on our behalf.

³ As I use 'knowing-wh' it refers to any form of knowledge which we attribute using knowledge ascriptions where the complement of 'knows' is an interrogative clause headed by a question word like 'when', 'why', 'whom', and 'how'. On this usage then knowing-how is a form of knowing-wh. Sometimes the term is used in a way that excludes knowing-how, but I take the inclusive usage to be more standard and helpful (Parent 2014: 81).

Hawley also suggests that the influence of the client's perspective might help to explain cases identified by Hawthorne (2000) and Stanley (2011) where people have the intuition that a subject possesses a given form of knowing-wh even though the subject's relevant true belief fails to meet standard anti-luck conditions on the possession of knowledge-that (we will discuss related cases in detail in §4). Hawley thinks that if I am interested, for example, in whether John knows where to find decent coffee because I want John to fetch a decent coffee for me, then "it is enough that John has a true belief about where to find decent coffee (and that he is willing to help)" (2011: 289).

Hawley's discussion of knowing-wh indicates how the client's perspective might be important, at least in certain contexts, to our thought and talk about knowledge-wh in general. In which case, we might conclude that, in these contexts, ascriptions of these other forms of knowing-wh are used to attribute forms of downstream rather than upstream knowledge. As with knowing-how, the idea will be that one has to be in a true belief state which not only has the right content (i.e., a proposition which answers the embedded wh-question) but which also explains why the subject in the state has certain associated abilities. So, for example, in the case of knowing when to invest in the stock market one will have a true belief about the right time to invest in the stock market and this true belief state will be possessed in such a way that it explains one's ability to invest in the stock market at the right time.

If there is a difference between knowing-how and these other forms of knowing-wh it is perhaps this: the downstream concept of knowledge seems to dominate our thought and talk about knowing-how, whereas with the other examples of knowing-wh that Hawley discusses it is only in quite specific contexts that our interest seems to be in downstream knowledge when we ascribe these forms of knowing-wh.

I am unsure though about the extent and significance of this disanalogy. For, on the one hand, the downstream concept of knowledge often seems to be in play when we ascribe forms of knowing-wh that are closely associated with certain abilities (like e.g., knowing where to find X and having an ability to locate X in some way, or knowing what to do and having an ability to do the correct thing). And, on the other hand, I think there will be some contexts where we use knowing-how ascriptions to attribute upstream rather than downstream knowledge, even though paradigmatically we use knowing-how ascriptions to attribute downstream knowledge (and I suspect the reverse is true for knowing-that ascriptions as some of Hawthorne's cases involve knowing-that rather than knowing-wh).

However, even allowing that there is a significant disanalogy here there are other examples where the downstream concept of knowledge and the client's perspective are, arguably, just as central to our thought and talk about a given form of knowledge as it is for knowing-how. Or, at least, that is what I will argue in the next two sections.

3. Insufficiency and the Client's Perspective

Many arguments for thinking that there are forms of knowledge other than knowing-that begin with certain intuitions that mere knowledge-that (often gained through testimony) is not sufficient for possessing the relevant form of knowledge. This is certainly the case for knowing-how,⁴ but insufficiency cases also play similar roles in debates about other important forms of knowledge and understanding. What I want to explore now is how the client's perspective, and the idea of downstream knowledge, can help to provide us with an attractive and unified explanation of these insufficiency cases, and to illuminate these different forms of knowledge. I will focus on the two case studies of understanding-why and 'what it is like'-knowledge, but I think many of the lessons here will generalise to other examples.

3.1 Understanding-why and insufficiency

Insufficiency cases concerning understanding-why help to motivate one of the most important views in debates about moral testimony. The view is that moral testimony can provide us with mere moral knowledge-that but not moral understanding-why, which is either equated with moral knowing-why or is thought to be distinct from both moral knowing-that and knowing-why. Hopkins (2007), for example, suggests that on the basis of the pure testimony of one's colleagues alone one could gain the moral knowledge *that* you should all strike, but one could not thereby come to know *why* you should strike. And Hills (2009) claims that based solely on a friend's testimony Eleanor could come to know that eating meat is wrong, and she could even come to know why it is wrong, but she could not come to *understand why* it is wrong on this basis.

Why isn't mere moral knowing-that sufficient for moral understanding-why? The usual suggestion is that understanding-why in general involves the possession of certain abilities which are not required to merely possess the relevant kind of knowing that. Hills (2009: 102),

⁴ For example, when Stanley and Williamson (2001: 428-429) introduced the notion of a 'practical mode of presentation' it was to try to address an insufficiency case for knowing-how involving someone who knows that *that* way is a way to ride a bicycle but, nonetheless, does not know to ride a bicycle.

for example, holds that if you understand why p (and q is why p), then you will have abilities (at least to some extent) to do (i)-(vi):

- (i) follow an explanation of why p given by someone else;
- (ii) explain why p in your own words;
- (iii) draw the conclusion that p (or that probably p) from the information that q ;
- (iv) draw the conclusion that p' (or that probably p') from the information that q' (where p' and q' are similar to but not identical to p and q);
- (v) given the information that p , give the right explanation, q ;
- (vi) given the information that p' , give the right explanation, q' .

To see how the client's perspective can apply to our thinking about moral understanding-why, imagine that Eleanor's friend, let's call her Hannah, ended up convincing Eleanor to join her as a fellow member of an animal rights group. A few months later, Hannah, is invited to an important debate with a representative of the meat industry on the ethics of eating meat, but Hannah is unable to attend. If Hannah is considering Eleanor as a possible replacement, then part of what she will want to assess is whether, and to what extent, Eleanor has abilities to perform the kinds of actions of making inferences and giving explanations etc. that Hills outlines in (i)-(vi).

Hannah, we can assume, has lots of relevant true beliefs about why eating meat is wrong, and she possesses this information in ways that explain her own abilities to perform these actions. But, in this situation, Hannah is a client seeking someone to perform actions of moral explanation on her behalf, rather than an inquirer or an apprentice seeking someone who can help her obtain knowledge for herself. Hannah's interest then in assessing whether Eleanor understands why eating meat is wrong, is an interest not just in whether Eleanor believes the right kinds of true propositions about this moral issue, but whether Eleanor thereby possesses the right kinds of abilities.

Now suppose that the function of our concept of understanding-why is to serve our interests as clients in identifying reliable performers of such actions of inference and explanation, and that, accordingly, understanding-why is a kind of downstream knowledge defined in terms of such abilities. These twin hypotheses are well-placed to explain insufficiency cases like Hills' Eleanor case.

For if we take understanding-why to be a kind of downstream knowledge the idea will be that understanding why p is a matter of having at least one true belief that q (for some proposition q that is an answer to the question ‘why p ?’) where this true belief is possessed in such way that it explains, and perhaps even entails, one’s having certain associated abilities. And if Hills is right these will be abilities to perform at least some of the actions in (i)-(vi). In which case, it is clear why mere *upstream* knowledge that q is not a sufficient condition by itself for understanding why p , because the kinds of etiological properties required for a true belief that q (which answers the embedded why-question) to constitute a state of upstream knowledge do not include the kinds of action-guiding properties required for it to constitute a state of downstream knowledge which can explain one’s having such abilities.

Of course, one could explain why mere knowledge-that is not sufficient for understanding-why without bringing in these broader ideas of the client’s perspective and downstream knowledge. But I think locating understanding-why within this framework can help us to get a clearer picture of the relationships between understanding-why and knowledge, both at a very general level and in relation to the details of existing debates about this relationship.

The large-scale picture is this: understanding-why is a species of downstream knowledge, and the function of our concept of understanding-why is something like identifying reliable explainers,⁵ which is a species of the more general function of our concept of knowledge, stemming from the client’s perspective, which is identifying reliable performers (see §5 for thoughts on how the function of identifying reliable informants fits into this picture). And one example of an existing debate that this framework can help us to clarify is debates about Hills’ claim that understanding why p cannot be identified with knowing why p .

Hills’ (2009) supports this claim by imagining that, based on Hannah’s pure moral testimony again, Eleanor comes to know the further proposition that eating meat is wrong *because* of the suffering of animals under modern farming methods (which is assumed to be a correct answer to the question ‘why is eating meat wrong?’). Hills assumes that Eleanor would thereby know why eating meat is wrong, but she would not thereby understand why eating meat is wrong because such understanding requires further abilities to do actions (i)-(vi) which are not

⁵ See Hannon (2019: Ch. 9) for a development of this idea. Hannon’s valuable discussion of understanding accords with much of what I argue here, and also what I argue below in §4.1 (see his 2019: 250-252). The main difference is that the framework I develop further illuminates the relationship between understanding and knowledge in the ways described above.

required to merely know that eating meat is wrong because of the suffering of animals under modern farming methods.

On the other hand, Hopkins (2007) seems to equate moral understanding-why with knowing-why in his strike example, and similarly Howell (2014: 397) claims in response to Hills that “while there is a clear difference between knowing p and understanding why p , it is not as clear that there is a difference between understanding why p and knowing why p ”. Howell does not dispute Hills’ claim that understanding-why involves certain abilities over mere knowing-that, but he thinks that “knowing why p would typically involve many of these abilities as well” (*ibid*: 397).

These conflicting intuitions about knowing-why can be nicely explained by the upstream versus downstream distinction, and three plausible ideas: (1) understanding-why ascriptions are paradigmatically used to attribute a form of downstream knowledge (supported by our discussion above in §3.1); (2.1) in some contexts knowing-why ascriptions are used to attribute the same form of downstream knowledge, but (2.2) in other contexts they are used to attribute upstream knowledge (supported by our discussion in §2); and (3) knowing-that ascriptions are paradigmatically used to attribute upstream knowledge (supported by orthodox assumptions about knowledge-that).

Hills assumes that knowing that q , for some proposition q which is a correct answer to the question ‘why p ’, is a sufficient condition for knowing why p . This assumption makes sense given (3) and if, following (2.2), we only have in mind those contexts where knowing-why ascriptions are used to attribute upstream knowledge. In which case, one could reasonably conclude, in line with (1), that understanding-why is distinct from knowing-why. But if, following (2.1), we focus on those contexts where knowing-why ascriptions are used to attribute downstream knowledge then the distinction between knowing-why and understanding-why will collapse in the sense that *in those contexts* a knowing-why ascription and the parallel understanding-why ascription will both attribute the same epistemic state. In which case, given (3), in such a context the relevant knowledge-that ascription will likely not entail the truth of either the knowing-why or the understanding-why ascription. I think this is the correct point that Howell is picking up on when he suggests that the ‘understanding strategy’ should be reformulated in terms of the claim that knowledge-that does not suffice for knowing-why (2014: 397).

On this view then there is a good sense then in which Hills and Howell are each partly right, and each partly wrong, on the issue of whether knowing-why is distinct from understanding-why. Hills is right, and Howell is wrong, in the sense that knowing-why ascriptions are sometimes used to attribute a different epistemic state from that attributed by understanding-why ascriptions. However, Howell is right, and Hills is wrong, in the sense that knowing-why ascriptions are sometimes used to attribute the same epistemic state as that attributed by understanding-why ascriptions.

3.2 'What it is like' knowledge and insufficiency

Consider the knowledge you have when you know *what it is like* to, say, fall in love, eat vegemite, see something red, or be a parent. Can such 'what it is like' ('WIL' for short) knowledge be identified with mere knowing-that which we can gain from sources like testimony or books? It is often assumed that it cannot (Lewis 1988, Paul 2014) on the grounds that WIL-knowledge is subject to an *experience condition* such that, at least in any normal circumstances, one can know what it is like to Φ only if one has Φ -ed oneself. And this idea is, of course, related to the powerful intuitions that drive not only Jackson's thought experiment of Mary in her black-and-white room, but also everyday intuitions like the thought that, say, one cannot know what it is like to be in love until one has been in love.

Partly in response to such intuitions some theorists have been driven to views on which WIL-knowledge is a wholly non-propositional form of knowledge. Perhaps most notably Lewis (1990) and Nemirow (1980) offered the *ability hypothesis* according to which knowing what it is like to Φ is identified with the possession of certain abilities—for Lewis the abilities to imagine, remember, and recognise experiences of Φ -ing, and for Nemirow just the imagination ability—as opposed to any kind of propositional knowledge.

There are many problems, however, with the ability hypothesis, and especially with the idea that WIL-knowledge is a wholly non-propositional form of knowledge (see e.g., the ten arguments in Lycan 1996: 92-95). For these reasons most theorists now accept that WIL-knowledge can be analysed in line with the standard question-answer analyses of knowing-wh, so that knowing what it is like to Φ is at least partly a matter of knowing some proposition p which is an answer to the embedded WIL-question (e.g., Hellie 2004, Lycan 1996, Stoljar 2016, Tye 2011). But it is also common to grant that merely knowing such a proposition isn't sufficient for knowing WIL to Φ because one also has to meet some further condition which could only be met if one had an experience of Φ -ing oneself (see e.g., Tye's 2011 view on

which one also has to possess an acquaintance-based concept of the way it feels to Φ). And, if knowing WIL to Φ is only partly a matter of knowing some proposition p which answers its embedded question, then there is a clear sense in which *mere* knowing-that is not sufficient for WIL-knowledge.

What light could the client's perspective shed on the insufficiency of mere knowing-that for possessing WIL-knowledge? In the case of understanding-why the support for the insufficiency claims comes from the assumption that this epistemic state involves the possession of *abilities*. But the insufficiency claims here seem to stem from WIL-knowledge requiring *experiences*. Doesn't this difference undermine the idea that the client's perspective is important to our thought and talk about WIL-knowledge?

The answer to such concerns can be found by reconsidering Lewis and Nemirow's ability hypothesis. It is true that most theorists today would reject all of the commitments of the view, and especially the claim that knowing-WIL is a wholly non-propositional form of knowledge. But that position is consistent with the plausible idea that there is still an intimate connection of some kind between knowing-WIL and the kinds of abilities that feature in the ability hypothesis, including even the claim that knowing-WIL entails having at least some of these abilities in normal circumstances. And it is these abilities which provide our entry point here for seeing how the client's perspective can be applied to WIL-knowledge.

Suppose you are chef who is badly in need of a holiday, and so you want to delegate your role to someone else. Many of your responsibilities that need replacing involve WIL-knowledge and its associated abilities, like, say, knowing what a dish tastes like when it has too much, too little, or just enough salt. Or when thinking of a new dish, you need to be able to imagine what it would taste like if you paired one ingredient with another. So, when you are considering who to put in charge you will have an interest in whether they possess information about what it is like to have various possible kinds of experiences, and that they possess this information in such a way that it supports their having abilities to imagine, recognise, and recall various kinds of taste and smell experiences. More generally, the client will have an interest in identifying people who do not merely possess correct information about what it is like to Φ but who possess that information in such a way that they can reliably perform related actions of correctly recognising experiences of Φ -ing when they have those experiences again, or accurately imagining or recalling such experiences.

But still if WIL-knowledge is a form of downstream knowledge why is it subject to an experience condition? The reason is that, as Lewis (1990: 78-79) discusses, while having an experience of Φ -ing is not the only possible way of coming to possess WIL-knowledge and its accompanying abilities (I will explore this claim in more detail in §4.2), it is the best and normal way of coming to possess that knowledge and those abilities. So, to possess, say, the ability to imagine experiences of Φ -ing, then, in normal circumstances, one must have had an experience of Φ -ing at some point oneself. In which case, the client's interest in identifying people with that ability will lead them to having an interest in identifying people who have had experiences of Φ -ing.

The idea that knowing-WIL is subject to an experience condition is consistent then with our hypotheses that the function of our concept of WIL-knowledge serves our interests as clients, and that WIL-knowledge is a kind of downstream knowledge. And these hypotheses can explain why mere knowledge of a proposition that answers the question 'what is it like to Φ ?' does not suffice for one to know what it is like to Φ , and why such WIL-knowledge cannot be gained solely from, say, consulting testimony about how experiences of Φ -ing feel. For while testimony can provide one with knowledge of propositions that address that question, it clearly is not a way to gain abilities to imagine, recognise, and recall experiences of Φ -ing.

Furthermore, as with understanding-why, there will be explanatory payoffs to situating WIL-knowledge within this larger framework. Consider those theorists who think of WIL-knowledge as a matter of knowing an answer to the WIL-question and meeting some further acquaintance or experience condition. One concern with such views is that it might seem ad hoc to just add an experience condition to the standard question-answer analysis of knowing-WIL. After all, why is this form of knowledge subject to an experience condition when propositional knowledge more generally is not subject to any such condition? Our discussion shows how one can explain why WIL-knowledge is subject to an experience condition as part of a larger framework that applies to our thought and talk about propositional knowledge in general.

4. Redundancy and the Client's Perspective

Arguments for thinking that there are important forms of knowledge or understanding other than knowing-that sometimes rely not on insufficiency claims, but on what I will call 'redundancy' claims. These are claims that possessing the relevant kind knowing-that is not

even *necessary* (and, hence, is redundant) for possessing the given form of knowledge or understanding.

While there is no consensus view of knowledge-that there is a lot of agreement that certain conditions are individually necessary for possessing knowledge-that; for example, that knowledge-that is respectively subject to an anti-luck condition, and a doxastic justification condition. In which case, if one can show that a given form of knowledge is not subject to one of these conditions then one can use that premise to support the conclusion that the given form of knowledge or understanding is not any form of knowledge-that.

Redundancy arguments against intellectualism have played a significant role in the knowing-how literature (e.g., Carter and Pritchard 2015, Cath 2011, Poston 2009). And closely related arguments have been made for an even longer time for thinking that various forms of understanding are distinct from knowing-that, including understanding-why both moral and non-moral (Hills 2009, 2015), understanding of a subject matter (Kvanvig 2003), and linguistic understanding (Pettit 2002).

What I want to show, first, in this section is how the client's perspective, and the idea of downstream knowledge, can help to explain what is going on in the redundancy cases involving understanding-why. I will then consider whether there might be similar redundancy cases for WIL-knowledge even though such cases have not been identified previously.

4.1 Understanding-why and redundancy

To see how redundancy claims can feature in arguments for the conclusion that understanding-why can come apart from both knowing-why and knowing-that, consider the following argument from Hills (2009: 104):

Suppose that your school has been sent a set of extremely inaccurate textbooks, which have been handed out to your class. But you are very lucky because there is only one that is accurate, and by chance you have it. You read in your book that Stalin was responsible for the deaths of millions of people. You draw the obvious conclusion that he was an evil person. It is plausible that you do not know that Stalin killed millions of people, since you could so easily have got a different textbook which was wholly unreliable. As a consequence, you do not know why he was evil (that is, you do not know that he was evil because he killed millions of people, because you do not know that he killed millions of people). But I think that you can understand why he was. After all, you believe that he was evil because he killed millions of people, and that is correct, and you have—let us assume—the ability to draw the conclusion that he was evil from the reasons why he was evil and to do the same in similar

cases. So it seems that you can have moral understanding why p without having knowledge why p .

The idea that this case is a counterexample to the idea that understanding-why can be identified with knowing-why is obviously based on two key claims, namely, that in this scenario: (1) one understands why Stalin was evil, but (2) one does not know why Stalin was evil.

Focusing on (1) to begin with, note that Hills supports (1) by pointing to the fact that in this scenario one does have a true belief about why Stalin was evil, and by making the stipulation that one would still possess the ability to infer that he was evil from the reasons why he was evil and to the same in similar cases. When Hills first introduces the abilities to do (i)-(vi) in her discussion the initial claim is that possessing some significant number of them is a necessary condition for understanding why p . But at this point in her discussion the idea seems to be that possessing the relevant true belief and some significant number of these abilities is a sufficient condition for understanding why p .

This claim fits well with the idea that understanding-why is a kind of downstream knowledge. For if the downstream hypothesis is correct then we would expect that someone who has a relevant true belief of the form ' p because q ', and who has abilities to do (i)-(vi), has all that they need to understand why p . For to say that understanding why p is a form of downstream knowledge is just to say that it is *solely* a matter of having a true belief with the right content and possessing that true belief in such a way that it explains one's having certain relevant abilities.

Similarly, if we think of the client's perspective, what they seek is someone who is in a true belief state that supports certain associated abilities, and *nothing more*. So, any further upstream properties of the relevant true belief state will be irrelevant to the client's search if those properties make no significant difference to whether the subject possesses these abilities. And, as Hills implicitly assumes, the kind of epistemic luck present in her Stalin case seems perfectly compatible with one's still having the kinds of explanatory and inferential abilities involved in possessing understanding-why.⁶

What about (2)? That is, Hills' claim that in her Stalin case one would not know why Stalin was evil. This claim rests on the orthodox assumption that all forms of knowing-that—

⁶ I agree with Hills' assumption here but see Pavese (2021) for important arguments that abilities to perform intentional actions require upstream knowledge that is incompatible with the epistemic luck found in Gettier-type cases. Ultimately, I don't think these arguments succeed but that discussion will have to be left for another occasion. For arguments that abilities to perform intentional actions only require downstream knowledge see Cath (2015a).

including any forms of knowing-wh, like knowing-why, that can be analysed in terms of knowing-that—are inconsistent with the kinds of epistemic luck present in Gettier-like cases (understood in a broad way to include cases involving environmental luck). In which case, if knowing-why is a kind of knowing-that, and understanding-why is present in the Stalin case, then this is an example where knowing why p comes apart from understanding why p .

But, drawing on our earlier discussion, we can see that this assumption is problematic given that, in some contexts, knowing-why ascriptions are plausibly used to attribute downstream knowledge. So, while examples like the Stalin case might provide us with good reasons to deny that understanding-why can be identified with *upstream* knowledge-why, they do not provide us with good reasons to deny that it can be identified with the *downstream* knowledge-why.

4.2 WIL-knowledge and redundancy

Redundancy claims have not featured in any notable way in discussions of WIL-knowledge, but if WIL-knowledge is a kind of downstream knowledge then we should expect to find such cases. And I think if we look more closely at existing discussions of WIL-knowledge we can see evidence of implicit commitments to the possibility of such cases.

As mentioned earlier, Lewis (1988: 78-79) held that while experience is the best way of acquiring WIL-knowledge it is not the only *possible* way. Lewis notes that there is some change that happens in a person (where he clearly has in mind an internal physical change) when they, say, smell a skunk for the first time and thereby come to know what it is like to smell a skunk. And then Lewis suggests that the exact same change could be caused to occur in someone by advanced neuroscience or magic, and that, if so, that person could know what it is like to smell a skunk without ever having smelled a skunk themselves.

These claims make sense for Lewis given his view that states of WIL-knowledge are just abilities, and it also make sense if we think of WIL-knowledge as a true belief state which supports the possession of such abilities, because neither kind of state requires the possession of etiological properties concerning the origins of these states.⁷ And we can view Lewis' (1988: 78-79) acknowledgement that some might “quibble about whether a state produced in this

⁷ More carefully, this is true when we restrict our attention to the abilities to imagine and recognise experiences of Φ -ing. The ability to remember experiences of Φ -ing clearly does require the upstream condition of having had an experience of Φ -ing at some point in the past. So, in the situations Lewis imagines the subject would only have an ability to have apparent memories of Φ -ing. Elsewhere Lewis (*ibid*: 98) discusses the possibility of losing the ability to remember but retaining the ability to imagine, and I suspect his view is that the later ability alone is sufficient for knowing what it is like to Φ .

artificial fashion would deserve the *name* ‘knowing what it is like to smell a skunk’” as a nod to the point that, like other forms of knowing-wh, knowing-WIL ascriptions can presumably be used in some contexts to attribute upstream knowledge (a point that is consistent with the assumption that the downstream use is more central).

If Lewis is on the right track here then there should be redundancy cases for WIL-knowledge that parallel the kinds of cases we find in the literature on understanding-why and knowing-how, and I think there are such cases. Consider Mary⁸ before she has ever left her black-and-white room. One day Mary is told that she will be given brief access to a previously locked cupboard in her room which contains hundreds of copies of a colour textbook which is full of colour patches with labels identifying every colour. Unbeknownst to Mary all but one of the hundreds of copies of the book have incorrect labels (e.g., the label under the green patch says ‘red’ and the label under the red patch says ‘green’, etc.). Luckily, Mary grabs the one copy with accurate labelling and then the cupboard is locked again. Looking now at the book Mary sees the red patch with the label ‘red’ underneath it and so she comes to truly believe that *this* is what it is like to see something red.

Does Mary now know what it is like to see something red? When the client’s perspective is salient, I suspect we will find the intuition that she does possess this knowledge (e.g., think of someone who is wondering if Mary could fetch them the red tube in an unlabelled box of paints). For Mary does now have a true belief state which explains her having abilities to imagine and recognise experiences of seeing something red, and if our interest is finding someone who can reliably perform such actions then we will only be interested in whether they are in such a state, not how they came to be in such a state.⁹

5. Knowledge Divided?

I have explored how Hawley’s notion of the client’s perspective, and the associated idea of downstream knowledge, can apply to not only knowing-how but also understanding-why and WIL-like knowledge. Obviously, these explorations are preliminary and a lot more could be said about these issues. But I hope to have shown how far one can go in explaining key features

⁸ To avoid any unnecessary commitments in relation to the knowledge argument I only have in mind ‘ordinary Mary’ here (Paul 2014: 9).

⁹ Currie (2020: 89) gives the only other ‘Gettier’ case for WIL-knowledge I am aware of, and his discussion accords well with my claims here as he also holds that the presence of Gettier-style luck is irrelevant to whether someone possesses WIL-knowledge.

of these two forms of knowledge within the framework of the client's perspective and downstream knowledge.

However, there are worries one might have with the idea of downstream knowledge which might lead some to quickly dismiss the value of our explorations here. In the knowing-how literature orthodox intellectualists often reply to related ideas by claiming that they should be rejected because: (i) they commit us to an ambiguity thesis about the meaning of knowing-wh and knowing-that ascriptions (with one distinct sense denoting upstream knowledge and the other denoting downstream knowledge), and (ii) linguistic tests provide us with good evidence that this ambiguity thesis is false (see e.g., Stanley 2011: 176). But there are at least three serious problems with this kind of reply.

One problem is that the relevant tests are good tests for strict ambiguities (where the distinct senses are unrelated), but not for polysemy (where the distinct senses are closely related); see Abbott (2013), and Sgaravatti and Zardini (2008) for discussion. And if 'knows' has two senses for downstream and upstream knowledge this would clearly be a case of polysemy rather than strict ambiguity (Cath 2015b).

Another problem is that even if we assume there is one univocal semantic value for 'knows' that linguistic assumption does not tell us what that semantic value is. Orthodox intellectualists like Stanley (2011: 180-181) presuppose that this semantic value will denote upstream knowledge, and that insofar as knowing-wh ascriptions are used in some contexts to pick out true beliefs states which are not states of upstream knowledge then those uses will merely reflect the pragmatics, rather than the semantics, of such ascriptions. But that assumption cannot be taken for granted here, as one could equally argue that the unique semantic value denotes downstream knowledge and that it is the upstream uses which merely reflect the pragmatics of knowledge ascriptions, or one could suggest that the one semantic value can denote both upstream and downstream knowledge in different contexts due to it having an indexical-like character, or due to 'modulation' (Antony 2001), or one of many other theoretical options (Cath 2015a, and 2015b).

Finally, suppose Stanley (2011) is right and that whenever a knowledge ascription is used to communicate that someone has downstream knowledge this use is only achieved through the pragmatics rather than the semantics of such ascriptions. Even so, given how frequently we seem to use knowledge ascriptions for this purpose, and how important such uses seem to be to how we think about many forms of knowledge, there would still be considerable reason for

us to be interested in gaining a clearer theoretical picture of these pragmatic uses of knowledge ascriptions and the states that we are tracking with them (Cath 2015b).

There is no easy way then of dismissing the idea of downstream knowledge by appeal to linguistic objections. However, there are related conceptual worries one might have about the upstream versus downstream distinction. For one might worry that if our thought and talk about knowledge is equalled pulled between two very different sets of conditions then this will force us to conclude that our concept of knowledge is inconsistent.

As Habgood-Coote (2019) discusses in relation to the inquirer/apprentice versus client perspectives this might turn out to be a pill we just have to swallow, and there are existing models of how to understand inconsistent concepts in epistemology. That said, I think we may be able to avoid this consequence. In an important precursor to Hawley's discussion, Moore (1997) suggests that situations in which one is "looking...just for someone who is reliable" (*ibid*: 174) may be more *basic* than the situations of the inquirer or apprentice where one is looking for a reliable informant or teacher. One way of unpacking and supporting Moore's thought is to notice that providing accurate testimony, or accurate instruction, or demonstrations of how to perform an action, are themselves just certain kinds of actions that people can perform (see also Habgood-Coote 2019: 9). In which case, our interest in identifying people who can reliably perform such actions can be seen as just a special case of our more general interest, as clients, in identifying people who can reliably perform actions for or with us. And upstream knowledge could be seen as just a special case of downstream knowledge but one which is distinguished by the fact that an ability to provide reliable information is an ability one can only possess only if one's relevant true beliefs satisfy certain upstream conditions.

I am hopeful then that the distinction between downstream versus upstream knowledge need not lead us to the view that our concept of knowledge is inconsistent in any deep way, although the issue obviously requires further investigation. My main concern here has just been to show the fruitfulness of expanding the client's perspective beyond the case of knowing-how. And, if anything, I think these expansions, reveal a more, and not less, unified picture of knowledge and understanding. For we have seen how a single theoretical framework can be applied not only to knowing-how, but understanding-why and WIL-knowledge, and that within that

framework we can offer a unified explanation of both the function of our thought and talk about these epistemic states, and the nature of these states themselves.¹⁰

References

- Abbott, B. (2013) 'Linguistic Solutions to Philosophical Problems: The Case of Knowing How', *Philosophical Perspectives* 27(1): 1–21.
- Antony, M. V. (2001) 'Is 'Consciousness' Ambiguous?', *Journal of Consciousness Studies* 8(2): 19–44.
- Carter, J. A. & Pritchard, D. (2015) 'Knowledge-how and Epistemic Value', *Australasian Journal of Philosophy*, 93: 799–816.
- Cath, Y. (forthcoming) 'Social Epistemology and Knowing-How', to appear in J. Lackey & A. McGlynn (eds.), *Oxford Handbook of Social Epistemology*, Oxford: OUP.
- Cath, Y. (2020) 'Know How and Skill: The Puzzles of Priority and Equivalence', in E. Fridland & C. Pavese (eds.), *Routledge Handbook of the Philosophy of Skill and Expertise*, New York: Routledge.
- Cath, Y. (2017) 'Intellectualism and Testimony. *Analysis*, 77 (2): 1–9.
- Cath, Y. (2015a) 'Revisionary intellectualism and Gettier', *Philosophical Studies*, 172 (1):7–27.
- Cath, Y. (2015b) 'Knowing How and 'Knowing How'', in C. Daly (ed.), *The Palgrave Handbook of Philosophical Methods*, 527–552. London: Palgrave-Macmillan.
- Cath, Y. (2011) 'Knowing How Without Knowing That', in J. Bengson & M. Moffett (eds.), *Knowing How: Essays on Knowledge, Mind, and Action*, Oxford: OUP.
- Craig, E. (1990) *Knowledge and the State of Nature*, Oxford: OUP.
- Currie, G. (2020) *Imagining and Knowing: The Shape of Fiction*, Oxford: OUP.
- Glick, E. (2012) 'Abilities and Know-How Attributions', in J. Brown & M. Gerken (eds.) *Knowledge Ascriptions*, Oxford: Oxford University Press.
- Habgood-Coote, J. (2019) 'What's the point of knowing how?', *European Journal of Philosophy*, 27 (3): 693–708.
- Hannon, M. (2019) *What's the Point of Knowledge? A Function-First Epistemology*, New York: OUP.
- Hawthorne, J. (2000) 'Implicit belief and a priori knowledge,' *Southern Journal of Philosophy*, 38 (Supplement): 191–210.
- Hawley, K. (2011) 'Knowing How and Epistemic Injustice', in J. Bengson & M. Moffett (eds.), *Knowing How: Essays on Knowledge, Mind, and Action*, 283–99. Oxford: OUP.

¹⁰ My thanks to two referees for their helpful feedback. I would also like to thank Katherine Hawley for her inspiring work, and for being a wonderful colleague to me when I was at St Andrews.

- Hawley, K. (2003) 'Success and Knowledge-How', *American Philosophical Quarterly*, 40 (1): 19–31.
- Hellie, B. (2004) 'Inexpressible Truths and the Allure of the Knowledge Argument', in Y. Nagasawa, P. Ludlow, and D. Stoljar (eds.) *There's Something About Mary: Essays on Phenomenal Consciousness and Frank Jackson's Knowledge Argument*, 333–64. Cambridge, MA: MIT Press.
- Hills, A. (2015) 'Understanding Why', *Noûs* 49 (2): 661–688.
- Hills, A. (2009) 'Moral Testimony and Moral Epistemology', *Ethics* 120 (1): 94–127.
- Hopkins, R. (2007) 'What is wrong with moral testimony?', *Philosophy and Phenomenological Research* 74 (3): 611–634.
- Hookway, C. (2006) 'Epistemology and Inquiry: The Primacy of Practice', in S. Hetherington (ed.), *Epistemology Futures*, 95–110. Oxford: Oxford University Press
- Howell, R. J. (2014) 'Google Morals, Virtue, and the Asymmetry of Deference', *Noûs* 48 (3): 38–415.
- Kvanvig, J. L. (2003) *The Value of Knowledge and the Pursuit of Understanding*, Cambridge: CUP.
- Levy, N. (2014) 'John Bengson and Marc A. Moffett, eds., *Knowing How: Essays on Knowledge, Mind, and Action*. Reviewed by', *Philosophy in Review* 34 (6):284–286.
- Lewis, D. (1988) 'What Experience Teaches', reprinted in Y. Nagasawa, P. Ludlow, and D. Stoljar (eds.) *There's Something About Mary: Essays on Phenomenal Consciousness and Frank Jackson's Knowledge Argument*, Cambridge, MA: MIT Press.
- Löwenstein, D. (2017) *Know-how as Competence. A Rylean Responsibilist Account*, Frankfurt am Main: Vittorio Klostermann.
- Lycan, W. G. (1996) *Consciousness and Experience*, Cambridge, MA: MIT.
- Moore, A. W. (1997) *Points of View*, Oxford: Clarendon Press.
- Nemirow, L. (1980) 'Review of *Mortal Questions* by Thomas Nagel', *Philosophical Review* 89: 473–77.
- Paul, L. A. (2014) *Transformative Experience*. Oxford: OUP.
- Parent, T. (2014) 'Knowing-Wh and Embedded Questions', *Philosophy Compass* 9 (2): 81–95.
- Pavese, C. (2021) 'Know-how, action, and luck', *Synthese*, 198 (7): 1595–1617.
- Pavese, C. (2015) 'Practical Senses', *Philosophers' Imprint*: 15 (29): 1–25.
- Pettit, D. (2002) 'Why knowledge is unnecessary for understanding language', *Mind* 111 (443): 519–550.
- Poston, T. (2009) 'Know How to Be Gettiered?' *Philosophy and Phenomenological Research* 79 (3): 743–747.
- Paul, L. A. (2014) *Transformative Experience*, Oxford: OUP.
- Snowdon, P. (2004) 'Knowing how and knowing that: A distinction reconsidered', *Proceedings of the Aristotelian Society* 104 (1): 1–29.
- Stanley, J. (2011) *Know How*, New York: OUP.

Stanley, J. and Williamson, T. (2001) 'Knowing How', *Journal of Philosophy* 98: 411–44.

Stoljar, D. (2016) 'The Semantics of 'What It's Like' and the Nature of Consciousness', *Mind* 125/500: 1161–98.

Tye, M. (2011) 'Knowing What It Is Like', in J. Bengson & M. Moffett (eds.), *Knowing How: Essays on Knowledge, Mind, and Action*, 300–13. Oxford: OUP.

Hickman, N. W. (2019) 'Knowing in the "Executive Way": Knowing How, Rules, Methods, Principles and Criteria', *Philosophy and Phenomenological Research* 99 (2): 311–335.