

Moral understanding and knowledge

Amber Riaz

Published online: 21 May 2014
© Springer Science+Business Media Dordrecht 2014

Abstract Moral understanding is a species of knowledge. Understanding why an action is wrong, for example, amounts to knowing why the action is wrong. The claim that moral understanding is immune to luck while moral knowledge is not does not withstand scrutiny; nor does the idea that there is something deep about understanding for there are different degrees of understanding. It is also mistaken to suppose that grasping is a distinct psychological state that accompanies understanding. To understand why something is the case is to grasp or see why something is the case and grasping and seeing are ways of knowing.

Keywords Understanding · Knowledge · Grasping · Gettier · Internalism

1.

This paper defends the following claim.

MUK

Moral understanding is a species of knowledge.

Consider the proposition, P, that that action is wrong. MUK is the view that understanding why P amounts to knowing why P. It equates understanding why P with, for some Q, knowing that (P because Q).¹ The equivalence between knowing

¹ I use the word “because” here broadly. The “because” in “P because Q” is not just a causal one; it covers a range of relations between different classes of facts: the relations of supervenience, being a reason for, and so on.

A. Riaz (✉)
Lahore University of Management Sciences, D.H.A, Lahore 54792, Pakistan
e-mail: amber@lums.edu.pk; amberrz@gmail.com

why P and knowing that P because Q shows that knowledge why is *propositional*. Maybe in some contexts understanding why P requires ‘knowing why why P’, that is, for some Q and R knowing that (P because (Q and ((P because Q) because R))). In other words, in some contexts, understanding why P requires not just knowing the reason but knowing why it is a reason; that requirement would still be perfectly consonant with MUK.

The following example illustrates these points. The example is taken from Jane Austen’s novel *Emma*. On one occasion Emma’s treatment of Miss Bates is less than morally appropriate. In the interest of making a witty remark at a picnic with a large gathering of people, Emma publicly mocks the talkative Miss Bates on the dullness of her comments without realizing the insolence of her remark. But Mr Knightley does, and when he has the opportunity, he hammers that point home to Emma. He not only tells her that her treatment of Miss Bates was wrong, but gives the following, extended explanation of it to an initially bewildered Emma.

Were she [Miss Bates] a woman of fortune, I would leave every harmless absurdity to take its chance, I would not quarrel with you for any liberties of manner. Were she your equal in situation— but, Emma, consider how far this is from being the case. She is poor; she has sunk from the comforts she was born to; and, if she lives to old age, must probably sink more. Her situation should secure your compassion. It was badly done, indeed! You, whom she had known from an infant, whom she had seen grow up from a period when her notice was an honour, to have you now, in thoughtless spirits, and the pride of the moment, laugh at her, humble her—and before her niece, too—and before others, many of whom (certainly some,) would be entirely guided by your treatment of her...²

Mr. Knightley has moral understanding that Emma has acted wrongly; he sees or grasps why Emma’s action is wrong. MUK claims that his moral understanding of why it is the case that Emma has acted wrongly consists in his knowing the facts that make it the case that Emma has acted wrongly. In some other contexts, such as a moral philosophy discussion in which Mr. Knightley is participating, he would be required to have an extent of theoretical understanding of why Emma’s action qualifies as being wrong. It is consonant with MUK that in the context of a moral philosophy discussion, Mr. Knightley would be required to know why the identified facts make it true that Emma has acted wrongly.

The following sections defend MUK primarily by undermining the key motivations for Not-MUK. Most of the little literature on the relation between understanding and knowledge concerns non-moral understanding and knowledge. So, while this paper is mainly about moral understanding and knowledge, occasionally I will need to discuss the view that understanding in general is a species of knowledge; I shall refer to the latter view as ‘UK’.

² Source of text: <http://www.gutenberg.org>.

2.

One reason given in support of the view that moral understanding and moral knowledge are distinct is that while knowledge has an anti-luck constraint on it, understanding does not. And if understanding is a species of knowledge, then it must have the properties of knowledge. The following cases may be provided in support of this claim.

Case 1

You read a book in which the author has got the facts about Stalin and his killings correct by chance.

It might be thought that while you do not know that Stalin was evil, and why he was evil, you do understand why he was so. For example, Hills (2010) writes:

I am inclined to ... say that even if your own textbook was extremely inaccurate, but that you read in it the truth that Stalin killed millions, you could as a result understand why he was an evil person, though clearly you could not know that he was evil or know on that basis. I am inclined to think that you have moral understanding because your belief that Stalin killed millions is true and you formed it reasonably (assuming that you had no reason to think that your textbook was generally inaccurate) and you have grasped correctly he was evil because he killed millions – and this all that is required for understanding (2010: 196).

Note that Hills only mentions Case 1 in passing in a footnote. Her main ammunition against MUK is supposed to be the following case.

Case 2

[Y]our 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.

Hills' judgement about Case 2 is similar to her judgement about Case 1.

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 evil. 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 do the same in similar cases. So it seems that you can have moral understanding why p without having knowledge why p (2010: 196).

Hills models Case 2 on a case presented by Pritchard in (2009); she also follows Pritchard in her judgement that there is no knowledge in the case. Pritchard's example is in turn inspired by Kvanvig (2003). Both these cases and the authors' assessment are described below.

Burning House

...imagine that on finding one's house in flames, one approaches someone who looks as if she is the fire officer in charge and asks her what the reason for the fire is....the apparent fire officer that one asks about the cause of the fire is indeed the fire officer, but that one could nevertheless have been easily wrong because there were other people in the vicinity dressed as fire officers - all going to the same fancy dress party, say - who one could very easily have asked and who would have given one a false answer... the epistemic luck at issue would prevent it [the cognitive achievement] from counting as knowledge. The critical question for us, however, is whether it is a case of understanding. I want to argue that it is... After all, the agent concerned has all the true beliefs required for understanding why his house burned down, and also acquired this understanding in the right fashion. (Pritchard 2009)

Comanche

Consider, say, someone's historical understanding of the Comanche dominance of the southern plains of North America from the late seventeenth until the late nineteenth centuries. Suppose that if you asked this person any question about this matter, she would answer correctly. Assume further that the person is answering from stored information; she is not guessing or making up answers, but is honestly averring what she confidently believes the truth to be. Such an ability is surely constitutive of understanding, and the experience of query and answer, if sustained for a long enough period of time, would generate convincing evidence that the person in question understood the phenomenon of Comanche dominance of the southern plains. But does she have knowledge? Ordinarily, yes; but it is not required. For, on the usual theories of knowledge, all those answers could be given from information possessed and still fail to be known to be true, because the answers might only be accidentally true. For example, most history books might have been mistaken, with only the correct ones being the sources of the understanding in question and with no basis in the subject for preferring the sources consulted over those ignored. Such a case fits the model of a standard type of case found in the Gettier literature (in particular, the fake barn case), where such accidentally true beliefs are not justified in the way needed for the beliefs to count as knowledge. (p. 196–7)

Let us call the kind of luck at work in Case 2 (and in Burning House, and Comanche) “environmental luck” and the kind at work in Case 1 “veritic luck” (Pritchard 2009). In addition to several “benign” varieties of luck that are compatible with knowledge, there are some others, according to Pritchard (2005, 2009, 2010), that are incompatible with knowledge. Veritic luck concerns the truth of the belief in question. A belief is veritically lucky if “luck intervenes between the agent and the fact in such a way that the agent's belief is true nonetheless” (Pritchard 2010: 51); the luck intervenes “betwixt belief and fact” as it were. For example, your true belief that there is a sheep in the field formed while looking at a rock shaped like a sheep and while the sheep itself is hidden from view is veritically lucky. Environmental luck is different; it is not the kind of luck that intervenes

between agent and belief; no luck intervenes betwixt belief and fact in the case of environmental luck. Instead, the general environment in which the belief is formed is inhospitable as in Goldman's fake barn country scenario. Henry is driving through a fake barn country, and forms the belief that there is a fake barn while looking at the only real barn there. There is a direct causal link between Henry's belief and the fact that there is a barn, but the presence of fake structures nearby makes it seem like it is just a matter of luck that the belief is true.

The distinction between veritic and environmental epistemic luck is relevant to the current discussion because it may be argued that understanding (as well as knowledge) is *vulnerable* to veritic luck, so Case 1 does not provide any support to Not-MUK. It would be interesting to see how many people share Hills's inclination to judge that there is understanding in Case 1—I, for one, don't. In fact, the inclination in question was missing for all the people that I tested Case 1 on. (My victims included my family and some patient friends; they are all non-philosophers and all native speakers of the language Urdu.) The cases I used were modelled after Case 1, but about different subject matters, depending on the interest of my subject. Most subjects denied that there was understanding in the case given to them; one person thought that there was knowledge as well as understanding. What everyone was unanimous about was that the idea that there is understanding but no knowledge didn't "even make sense". A plausible explanation of this is the interesting fact that there is more than one word in Urdu that translates into 'knowledge', and in some contexts, 'understanding' and 'knowledge' are used interchangeably. With the help of some bilingual acquaintances who know both English and Urdu, I was able to identify about five different words for the English word 'knowledge': 'patta'; 'janana'; 'ilm'; 'samajh'; 'hikmat'. What's most interesting is that the word 'samajh' was the translation of both 'understanding (why)' and 'knowledge (why)', and I couldn't find any contexts in which they came apart, i.e., contexts in which 'samajh' was the translation of either 'understanding' or 'knowledge', but not both. It is interesting that the judgement of my subjects didn't change when I revised the cases in the light of people's initial comments to ones where the subject is repeatedly in Case 1-type scenario, so that most of her moral training³—in some sense of the word 'training'—is based on exposure to cases where veritic luck is at play. This judgement appealed to the distinction between "being successful", and "knowing and understanding". In particular, they said that the subject in cases like Case 1 may go through life being *successful* with respect to the particular domain in which they've had this pseudo training, but he or she wouldn't have the relevant knowledge and understanding.⁴

These examples are particularly important in our assessment of Hills's view. Hills is trying to motivate a mainly internalist theory of understanding with the help of Case 1. But at least one thing my small survey described above shows is that

³ Or whatever non-moral training was relevant in a particular case.

⁴ In fact, for what it's worth, some people pressed for more details of the case and even offered their diagnosis of why it might seem to someone that there is understanding in Case 1.

there is a range of intuitions about cases modeled on Case 1, and it is at least very questionable that the one that Hills needs is the most common one.

In general, it is not plausible that there is understanding in Case 1. However, to support Not-MUK, it would be sufficient to show that moral knowledge and moral understanding come apart in any *one* of the two cases above, Case 1 and Case 2. To this end, the defender of Not-MUK could claim that moral understanding might be vulnerable to veritic luck, but unlike moral knowledge, it is immune to *environmental luck*.⁵ In fact, this is exactly like the line of argument adopted by Duncan Pritchard in (2009). Arguing against Kvanvig's claim in (2003) that understanding is immune to all kinds of epistemic luck, Pritchard claims that both knowledge and understanding are incompatible with veritic luck, so they (knowledge and understanding) do not come apart in cases of veritic luck. However—Pritchard claims—knowledge is, and understanding is not, incompatible with environmental luck. Therefore, examples in which environmental luck is in play *would* give Kvanvig what he needs in support of his view: a kind of case where knowledge and understanding come apart. In particular, while Case 2 involves environmental luck and not veritic luck, it might still be used in support of Not-MUK. For, it might be argued on Hills' behalf that the environmental luck in play in it is compatible with moral understanding, but not with moral knowledge.

However, once the distinction between different kinds of luck is made and accepted, it becomes questionable whether knowledge also is vulnerable to both kinds of luck. Clearly, there's no convergence of judgement about this among non-philosophers and philosophers. I myself found Hills' verdict about Case 2 baffling when I first read it. *All* the non-philosophers that I spoke to regarding Case 1 also thought that there is knowledge (and understanding) in Case 2-type scenarios. A couple of people who went on to defend their response suggested that the view that there is no knowledge in Case 2 threatens to take away a lot of our knowledge. It is also interesting to note that the first reaction of all my Pakistani undergraduate students so far to the paradigm example of environmental luck, i.e. the fake barn case, has been that there *is* knowledge in it.⁶ Besides, even some philosophers maintain that knowledge is *compatible* with environmental luck.⁷ For example, Hawthorne (2003) writes:

I give six children six books and ask them each to pick one of the books at random. All but one contains misinformation about the capital of Austria. I ask the children to look up what the capital of Austria is and commit the answer to memory. One child learns 'Belgrade,' another 'Lisbon,' another 'Vienna,' and so on. I ask an onlooker who has witnessed the whole sequence of events (or someone to whom the sequence of events has been described) 'Which one of the schoolchildren knows what the capital of Austria is?' or 'How many of the children know what the capital of Austria is?' It is my experience that those

⁵ Pritchard (2009).

⁶ Whether or not Burning House, Comanche and Case 2 are in fact exactly like the fake barn case is an interesting question in its own right.

⁷ See Grimm (2006), Hawthorne (2003).

presented with this kind of case will answer, not by saying ‘None of them,’ but by selecting the child whose book read ‘Vienna’—even though that child was only given the correct answer by luck. (Hawthorne [2003], pp. 68–9)

I will assume that the book with the correct answer had ‘Vienna’ as a result of an appropriate causal chain, not a random process, otherwise the case involves veritic rather than environmental luck and the knowledge verdict is quite implausible. Agreeing with Hawthorne, Grimm (2006: 529) claims that “... someone can come to *understand* various things about the Comanches on the basis of good information sources in a bad information environment ... in cases of this sort there is also a strong tendency to say that the person *knows* these things about the Comanches as well.” If these philosophers are right, then there isn’t a plausible case for Not-MUK. For, if there is *neither* understanding *nor* knowledge in cases of veritic luck, and there is *both* knowledge *and* understanding in cases of environmental luck, then knowledge and understanding sway together, a result that does not further the case of Not-MUK: in Case 2, in addition to understanding why Stalin is evil, we arguably also *know* that Stalin is evil. So contra Hills, the example does not show that moral knowledge and understanding come apart in cases involving environmental luck.

Actually, Hills does not offer any reason to think that there is no moral knowledge in Case 2. She only acknowledges that her example is modeled on Burning House. Pritchard thinks there is understanding in Burning House, but no knowledge; presumably, the absence of knowledge is to be explained with reference to *safety*. In Pritchard (2010), the safety principle is explained as follows.

The Safety Principle

S’s belief is safe iff in most near-by possible worlds in which S continues to form her belief about the target proposition in the same way as in the actual world, and in all very close near-by possible worlds in which S continues to form her belief about the target proposition in the same way as the actual world, her belief continues to be true.

Knowledge, on Pritchard’s view, requires safety. Your belief in Burning House is unsafe hence it doesn’t qualify as knowledge. On an analogous diagnosis of Case 2, there is no knowledge in it as the belief that Stalin is evil because Q is unsafe. So let the target proposition be

T

Stalin is evil because he killed millions of innocent people unnecessarily.

Belief in T is unsafe if one or more of the conditions given in the Safety Principle are not satisfied by it. Let us suppose there is an uncontroversial characterisation of “the same way”, such as reading a history book.⁸ Then on the most natural construal, belief in T is unsafe in the sense that there is a close possible world where

⁸ It is questionable whether reading T in a different book that is also so very different from the book read in the actual world counts as the “same way or method”. Is the way in question *reading this particular book* or *reading any book*, for example? In the interest of keeping things simple, I leave this issue aside and assume the safety-based account has a more or less satisfactory answer to this question.

you pick up an inaccurate book that gives a different and incorrect reason for the claim that Stalin is evil. For example, it says (R) ‘Stalin is evil because he killed all and only new born babies for fun’. Maybe some books say (V) ‘Stalin was a good person because he was the saviour of humanity.’ (One can imagine many other claims like R and V.) Is your belief that T unsafe in that in a close world you falsely believe R or V? The answer is not straightforward: R and V are so different from T that it is puzzling how belief in them in close worlds is supposed to have a bearing on the safety of belief in T at the actual world.⁹ It may be claimed that R and V are *similar enough* to T, in that all three concern Stalin’s being evil. However, *are* they similar enough to be relevant for the assessment of the safety of the belief that T at the actual world? Where does one draw the line for ‘similar enough’ to apply? It is arguable that we can only draw the line after *first* judging whether or not there is knowledge in the case under consideration; the judgement that a belief is safe or unsafe is itself then dependent on the judgement that the belief amounts to knowledge. If so, then we cannot justifiably ignore the (not uncommon?) intuition that there is moral knowledge in Case 2.

Of course, Hills is under no obligation to favour Pritchard’s theory of knowledge, or his account of safety, and other philosophers may appeal to completely different theories in support of the intuition that there is no knowledge in Case 2. However, it is a worthwhile thought that just the way a theory that grants knowledge in the standard Gettier-type cases makes knowledge too easy, a theory that denies knowledge in Case 2-type cases makes knowledge too hard. If that’s correct, then, it may be argued, the data that a theory of knowledge should respect as a starting point should include the judgement that there is knowledge in Case 2. Furthermore, even if a theory denies knowledge in Case 2, it is questionable whether it will have the resources to say that there *is* understanding in Case 2. Grimm (2006) forcefully argues that properly fleshed-out cases of environmental luck about which there is an inclination to say there is no knowledge are also the ones where there is an inclination to deny that there is any understanding.

An important lesson of the above discussion is the following. Case 2, which is Hills’s primary motivation for rejecting MUK, is like a duck/rabbit picture. We are uncertain about which it is, and can switch between seeing it as a good (non-Gettier) case and as a bad (Gettier) case. On the first way of seeing it, the case can be seen as an example of *both* knowledge that Stalin is evil *and* understanding why Stalin is evil. On the second way of seeing it, the case can be seen an example of *neither* knowledge that Stalin is evil *nor* understanding why Stalin is evil. Actually, each way of seeing it is compatible with MUK and a decision between the two ways is not needed to defend MUK, though I have tried to motivate the view that it is a good case. What needs to be highlighted is that it is illegitimate to switch between the two ways of seeing the case half way through, and decide the question of knowledge on one way of seeing and the question of understanding on the other way of seeing—which is what Hills does and so makes it look like a counterexample to MUK. It is

⁹ Recall that in at least some versions of the fake barn case, which is a classic example of environment luck, the proposition believed in close worlds remains the same, i.e., that there is a barn.

like claiming that the duck-rabbit is a beakless duck; that's an easy victory, and more importantly, only an apparent one.

2.

On a not uncommon view of moral understanding, it is “deep” and “rich”, deeper and richer than knowledge. This depth and richness can be unpacked in terms of the following putatively necessary (and possibly) sufficient conditions on moral understanding.

“Moral understanding involves a grasp of the relation between moral propositions and the reasons why they are true ... The grasp of the reasons that is essential to understanding involves a number of abilities ... If you understand why p (and q is why p), then in the right sort of circumstances you can successfully:

- (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);¹⁰
- (v) given the information that p, give the right explanation q;
- (vi) given the information that p', give the right explanation q'” (Hills 2010: 194).

Hills suggests that you can know why P without having the abilities (i) to (vi) and “that is one reason why understanding why P is not the same as knowing why P” (Hills 2010: 195).

Below, I argue that Hills' view fails to account for low degrees of understanding. In response, Hills could adjust her view in light of that concern, but insist that (i) to (vi) hold for very high degrees of understanding. She could then object that MUK cannot accommodate (i) to (vi) even for very high degrees of understanding partly because it cannot even allow for degrees of understanding. I show how this objection doesn't withstand scrutiny, and that although MUK is compatible with (i) to (v) for high degrees of understanding, the requirements are mistaken to begin with even for those high degrees.

First, note that understanding is not an all or nothing affair; instead, it comes in degrees. Granting that thought, Hills claims that each degree of understanding, corresponds to the degree to which each of the conditions (i) to (vi) is met.¹¹ So, on her view one would not qualify as understanding why P if any of the six conditions is not met at all. However, this requirement as a condition for *all* rather than some degrees of understanding is extremely stringent and implausible. To get to a sophisticated degree of moral or non-moral understanding, we need to build on lower degrees of it, starting with cases where there is very minimal understanding such that many of (i) to (vi) are not satisfied. Consider how someone reading a philosophy paper for the first time may understand why Anselm's ontological

¹⁰ Grimm (2006), though extremely sympathetic to the view that understanding is a species of knowledge, makes a similar claim about what understanding requires. Following James Woodward (2003), he speaks of understanding in terms of having an ability to answer “what-if-things-had-been-different?” questions.

¹¹ Hills (2010: 195).

argument doesn't work and yet not get why Descartes' doesn't—until they figure out and hence know why replacing “greatest conceivable being” with “most perfect being” doesn't render one version more robust than the other—yet it is plausible this student has *some* understanding of why some types of ontological arguments fail, though she fails many of (i) to (vi). Putting this in terms of UK, a less than fully sophisticated degree of understanding can mean either that one knows only some of the total set of propositions that jointly explain P, or that one is not able to extrapolate from the one case and make judgements about similar cases, and so on.

How, it may now be objected, can MUK accommodate the idea of degrees of understanding? For recall that MUK is the view that understanding why P is knowing why P, and knowing why P amounts to, for some Q, the propositional knowledge that (P because Q). Now, there are strong arguments by Jason Stanley (2007) against the view that propositional knowledge ascriptions are gradable; how can MUK be correct if ascriptions of understanding come in varying degrees of strength, but knowledge ascriptions do not?

This possible objection to MUK may be averted if MUK is clarified further. For MUK claims that 'S understands why P' is equivalent to 'S knows why P' and in a given context, that is equivalent to, for example: for some Q, S knows that (P because Q). MUK is committed to the claim that there are degrees to which S knows why P, not that there are degrees to which S knows that (P because Q). Every degree of knowledge why corresponds to some propositional knowledge, and that is to be understood along the following lines. Let us assume a coarse scale of measurement for degrees of knowing why P which specifies those degrees as 1, 2, 3 and so on. Then, for example, *one way* to represent the propositional knowledge corresponding to different degrees of knowledge why P is as given in UCK below. (The colon symbol ':' represents 'corresponds to'.)

UCK

Knowledge why P to degree 1: knowledge that (P because Q).

Knowledge why P to degree 2: knowledge that (P because (Q and ((P because Q) because R)))).

Knowledge why P to degree 3: knowledge that (P because (Q and ((P because Q) because R) and (((P because Q) because R) because S))).

.

.

...

The target of Stanley's arguments would be the items that appear to the right of the colon. For example, the knowledge that (P because Q), on his view, is not gradable because it lacks the two key features of gradable expressions. First, the occurrences of 'very well', 'very much', 'really' and so on in propositional knowledge ascriptions are not genuine semantic modifiers of the knowing relation; instead, they are pragmatic indicators. The evidence Stanley provides in support of this view is that negations of sentences such as 'John knows very well that Bush is the president' are unacceptable. In contrast, the negation of a sentence featuring a gradable expression with a semantic modifier—resulting in a sentence such as 'I do not like

Bill very much’—is perfectly natural. Furthermore, ‘know’ does not combine well with a modifier in a non-assertoric speech act. For example, ‘Do you know very well that Bush is the president?’ is very unnatural, while ‘Do you like Bush very much?’ is clearly not.

The second feature of a gradable expression is that it lends itself to natural comparative constructions, such as ‘better than’, but “‘better than’ is not a natural way to express comparisons between levels of epistemic position with ‘know’.” Stanley’s evidence for this is linguistic data; in addition, he offers alternative explanations of assertions featuring comparative constructions of ‘knows’ to the explanation that ‘knows’ is gradable.

But note that none of these points holds with respect to knowing why P. ‘John knows very well why Bush is the president,’ ‘John does not know very well why Bush is the president,’ ‘Does John know very well why Bush is the president?’ and ‘John knows better than Mary why Bush is the president’ are perfectly natural assertoric and non-assertoric speech acts. This is not because attributions of knowledge why pose a challenge to Stanley’s view, or that MUK is false because Stanley’s view is correct. Instead, what they highlight is that if we understand MUK properly, that is, in accordance with UCK, then MUK will not be vulnerable to Stanley’s arguments. What is to be born in mind is that MUK maintains that degrees of knowing why P correspond to knowing *more informative truths* of the form ‘P because Q’, i.e. knowing *fuller* answers to the question ‘Why P’. This amounts to knowing fuller explanations of P. So what is being graded, according to MUK, is the *content* of the knowledge, *not* the propositional knowing relation itself.

So far, I have argued that there are degrees of understanding, that MUK can allow for that fact, and that conditions (i) to (vi) do not hold for all degrees. But do (i) to (vi) hold for very high degrees of understanding, and if so, can MUK accommodate them? Suppose that (i) to (vi) do hold for very high degrees. Then MUK can allow for those abilities in terms of what may be reasonably expected of someone who knows fuller answers to the question ‘Why P?’. MUK can allow that someone who *really* or *fully* understands why P has a kind of hyper-knowledge involving, among other things, knowing why P, why P’, why P* (where the situations corresponding to the judgments that P, that P’ and that P* have slight differences), knowing why the differences between those situations warrant different judgements and so on. MUK can allow for this because the fuller explanations grasped by someone who really knows why p—knowledge that (P because (Q and ((P because Q) because R) and (((P because Q) because R) because S)))—will need to allude to what happens in close cases.¹² This hyper-knowledge can be sophisticated enough to result in, and account for, the satisfaction of (i) to (vi), and we cannot just stipulate – as Hills does – that someone can have this knowledge and still lack understanding.

The general point so far is that Hills’ conception of moral understanding and knowledge that motivates the rejection of MUK is seriously mistaken because it relies on a distorted picture of *degrees* of understanding and degrees of knowledge

¹² This moral epistemology can be developed in a way that is consistent both with Moral Generalism and with Moral Particularism.

why. But if this conception were revised to make it more plausible, then understanding why and knowledge-why would become indistinguishable.

However, whether or not MUK can accommodate the kind of features that Hills takes understanding why to have, (i) to (vi) as requirements *even* on very high degree of understanding are extremely questionable to begin with. In particular, Hills' view on what counts as a satisfactory explanation for condition (i), and "the right explanation" in (v) and (vi) is unreasonably restrictive. In the many examples Hills provides in her (2009) and (2010), the explanations provided by subjects who qualify as having moral understanding are the precise reasons that make a certain moral judgement correct;¹³ their understanding is so open to reflection that they can extract those reasons whenever needed. However, imagine a subject, Tara, who has a very high degree of moral understanding but much of the explanation she offers in support of a moral proposition is often significantly irrelevant to the case under consideration. As an avid reader of fiction and non-fiction from a young age, a keen watcher of films and a good observer of real life situations in which she has been disposed to advise friends and family when consulted for help, Tara has developed a very good understanding of various aspects of human situations, including the moral ones. However, Tara is a member of a deeply religious society, and since she has always had everyone's appreciation there, she is not in the least antagonistic or indifferent towards it. Instead, unbeknownst to her, she values it to the extent that she inadvertently couches her views on most matters, moral or otherwise, in terms congruous with those of her fellow folks. (If psychologists are to be believed, such self-deception along with self-deception of many other kinds is not only possible, but almost universally present.) When asked to explain a moral judgement, Tara will link the relevant reasons with the oracle or the will of the deity in a convoluted way such that while we could say that Tara can give *an* explanation, we could not justifiably say that she can give the "right" one as required by conditions (ii), (v) and (vi).

One motivation for the kind of "right explanation" required by (v) and (vi) could be some version of *internalism* about understanding, that is, the view that one always has reflective access to the grounds for one's understanding. Given this access, one could sift through the net of one's beliefs to pick out those that capture the "right explanation." But the case of Tara just described challenges this idea. Also, according to most philosophers including Hills, understanding is supposed to be *factive* and *non-transparent*. It is factive in the sense that you cannot understand why P if your relevant beliefs about why P are false. If you believe that Hitler was evil because the oracle says so, then you fail to understand why Hitler was evil, though you might think that you have that understanding. Given the factivity of understanding, its non-transparency doesn't come as a surprise. It may seem to me that I understand why P but I don't actually understand, since the facts are not as I take them to be; and it may seem to me that I do not understand why P when I

¹³ See, for example, Hills (2010: 190–93). On page 192, she explains in a lengthy footnote that the relevant reasons are those that make the believed moral proposition true and not "non-explanatory" evidence like testimony.

actually do. The point is that the features of factivity and non-transparency lend most natural support to externalism rather than internalism about understanding.

Finally, the restrictions on what counts as the right explanation on Hills' view unfairly exclude some intuitively legitimate explanations as they fail to satisfy (i) to (vi), while MUK can easily accommodate those explanations as plausible answers to 'Why P?' provided by someone with even a high degree of understanding. Suppose P is the proposition that that action is wrong. Then the kind of explanations I have in mind are: 'This action is just like that other one in which S had clearly acted wrongly;' 'His [the agent's] account of the situation doesn't really gel;' 'Something is fishy here,' and so on. There's evidence to suggest that even some non-moral experts are sometimes unable to articulate reasons for their (correct) judgements, and surely, if anyone understands why P, an expert does.¹⁴ Call such experts 'silent experts.' If there are moral experts, as I think there are, there is no good reason why they couldn't be such silent ones. Maybe those experts could *eventually* provide certain reasons. However, there are at least two issues that silent experts bring out. First, different explanations may be offered in response to 'Why P?' depending on various contextual factors, such as the aims and nature of the listener/audience, the available time, and so on. In particular, no reasonable case can be made for saying that any of these explanations has to be the canonical one. But shouldn't there be a canonical one if there is to be "the right explanation" as required by (v) and (vi)? Second, Hills implicitly assumes that "the right explanation" is more or less theoretical in nature—consider, for example, all the subjects in her examples who qualify as having moral understanding—but understanding doesn't have to be theoretical. While Tara has read a lot of books, let's consider Sara who has simply had lots of experience dealing with people, listening to their personal narratives, advising them about personal matters, and observing the results of her advice. Although Sara's judgement is nuanced, she is not very articulate, partly because she is not well read, and doesn't have much formal education. She's a paradigm example of a silent expert, who will often only have the most minimal explanations to offer in support of her judgements that we know to be mostly correct. It won't do to object that Sara is an expert in virtue of mere true belief, or even basic

¹⁴ In the book *Blink*, which is about the importance of "snap judgements", Malcolm Gladwell gives examples of different experts who made certain correct judgements in their domain without being able to fully articulate at the time why they thought they were correct. For example, the book begins and ends with the case of the identification of a kouros for the J. Paul Getty Museum in California. Evelyn Harrison, who was one of the world's foremost experts on Greek sculpture, was one of those consulted in the process. Her reaction on being presented *with* the kouros, and being told that the museum was about to purchase it after extended examination by various people, was: "I am sorry to hear that". But Harrison wasn't able to say what it was about the kouros that had promoted her reaction; she only had a sense that there was something amiss about it. When the art historian Federico Zeri saw the kouros, he found himself staring at the sculpture's fingernails. "In a way he couldn't immediately articulate, they seemed wrong to him" (Gladwell 2006: 5). Upon seeing the kouros, the first word that occurred to Thomas Hoving, the former director of the Metropolitan Museum of Art in New York, was 'fresh', and "fresh" is not supposed to be the right reaction to have to a two-thousand-year-old statue. Gladwell gives other examples of experts, ranging from those of marriage, war strategy, music, and so on, where the expert in question initially had the bare minimum explanation of their judgement to offer. For instance, a common kind of explanation they were able to think of was comparative claims, such as 'this is not like that', or 'there's just something not right about it'.

propositional knowledge, such as that that action is wrong, rather than understanding why or knowledge why, as an anonymous reviewer has suggested. The key motivation for this suggestion is that if one understands, one has the ability to offer a specific kind of explanation, i.e. a broadly theoretical one. It has already been argued that this motivation is ill founded. In addition, that she has understanding why (knowledge why) is a much better explanation of why she has a track record of giving good advice that works rather than the explanation that she has merely true belief or knowledge of the simple proposition that that action is wrong. Of course, all that requires optimism about the assumption that there are moral experts whose moral success is at least sometimes verifiable.¹⁵

3.

In ‘Is Understanding a Species of Knowledge?’ Stephen Grimm considers the following objection to the view that understanding is a species of knowledge. The psychological component of knowledge is generally taken to be belief. An analogue of belief is the act of assent: to believe that P is to assent to P, or to say ‘Yes’ to the claim that P. If knowledge is a species of belief, and understanding is a species of knowledge, then understanding is a species of belief. Given all this, it follows that the psychological component of understanding—that is, belief—is to be characterized along the lines of an act of assent. But what would be the claim to which one would be assenting in the case of understanding? When you believe that P, you would assent to the claim that P; when you understand why P, the candidate proposition that has a claim on your assent is that P because Q. However, what is supposed to be problematic about this suggestion is that an act of assent seems too “thin to capture what is going on when we take ourselves to understand” (Grimm 2006: 531). Grimm’s first example to explain this point concerns his Volkswagen whose gauges all seem dead 1 day. The mechanic explains that the gauges are dead because of a faulty ignition switch. Grimm then has “excellent reason to assent to the claim that [his car’s] gauges are dead because of a bad ignition switch, even though [he] fail[s] to grasp how a bad ignition switch might lead to this result.” His second example concerns mathematics. One can assent to the fact that a theorem is explained by a proof on the basis of one’s teacher’s testimony, but one may still fail to understand how the one explains the other. “Grasping the way in which the theorem depends on the elements you cite in your proof is different from assenting to the claim that the dependency holds” (Grimm 2006: 532).

The gist of the complaint being made is that the psychological component of understanding is grasping, and belief/assent or even piles of belief/assent simply do not amount to grasping; belief/assent is too impoverished to capture grasping. But this is supposed to be troubling for the view that understanding is a species of knowledge, as it is belief/assent that goes with the latter, not grasping.

Grimm’s own solution to the putative problem appeals to the idea that “while belief can be quite a thin psychological state, some forms of belief are thicker than others” such that “within belief there are simple acts of assent and then there are

¹⁵ I explore the notion of moral expertise and success in a separate paper.

also acts of assent that ... combine an element of grasping with the act of assent” (2006: 533). He demonstrates this point with the help of the example of “a priori knowledge”. What makes the proposition that $2 + 2 = 4$ an instance of a priori knowledge is not just its source for the source can be testimony for a subject S. But what does make it an instance of a priori knowledge for S depends on the nature of the psychological attitude S bears to the content mixed with the attitude of assent. Grimm suggests that this attitude is grasping.

I am not going to consider, or at least not in any detail, the merits and demerits of Grimm’s response. Apart from the fact that it makes reference to controversial notions, such as ‘a priori knowledge’, its suggestion that grasping is an independent psychological state is unconvincing. The status he accords to grasping is reminiscent of Descartes’s idea of rational intuition as a faculty that enables us to grasp certain truths though reflection, and I am not convinced that we need to mystify ‘grasping’ to solve what appears to be a pseudo-problem for MUK in the first place.¹⁶ In fact, the apparent problem disappears once one attends to the fact that what one knows comes in degrees of informativeness. Consider the above example about Grimm’s Volkswagen. It is very plausible that Grimm does know that his car’s gauges are dead because of a bad ignition; it’s just that he doesn’t know much more than this. The belief that his knowing involves is, therefore, going to appear to be “thin”. The mechanic, on the other hand, knows why the gauges are dead to a greater extent than Grimm for, he also knows why it is that a bad ignition causes the gauges of a car to be dead. While Grimm is only in a position to assent to the proposition that the gauges are dead because the ignition is faulty, the mechanic is in a position to assent to many more propositions, for example a proposition that explains why it is that gauges are dead because of a faulty ignition. Similarly, in the example about mathematics, the student has very little knowledge of the subject matter which explains the “thinness” of the belief it entails. At some very high level of one’s mathematical skill, one’s knowledge enables one to solve unfamiliar mathematical problems without any help, for example. At this high level, one would assent to many relevant proposition, for example, those featuring in a proof, the proposition that a theorem is explained by things cited in a proof, and so on. This is not to suggest that several assents “add up to a genuine grasping” in the sense of being sufficient for grasping—for you can keep assenting to various mathematical truths on the basis of testimony, and without having much knowledge of mathematics—but knowing lots will entail piles of assent, or rather piles of *whole-hearted* assents. For some, it may be tempting to think of grasping as a unique psychological state that accompanies understanding. But this temptation is easily resisted if we note that to grasp why something is the case is to understand why it is the case. And understanding why is equivalent to knowing why, knowing why is equivalent to knowing propositions of the form ‘P because Q’, and (on at least most plausible theories of knowledge) knowledge entails belief, so grasping why also entails belief.

¹⁶ In a recent paper that I discovered after having written this section, Pritchard voices a similar concern about Grimm’s views on grasping.

1 Concluding remarks

I have argued that moral understanding is a species of moral knowledge. Understanding why an action is wrong, for example, amounts to knowing why the action is wrong. Resistance to this view is often motivated by the implausible concern that moral understanding is immune to luck while moral knowledge is not. One other reason MUK is rejected is that moral understanding is thought to be deeper than knowledge. However, this idea does not withstand scrutiny as different levels of understanding correspond to different extents of knowledge. It is also mistaken to suppose against MUK that grasping is a distinct psychological state that accompanies understanding. To understand why something is the case is to grasp or see why something is the case and grasping and seeing are ways of knowing.

Acknowledgments I would like to thank Timothy Williamson, Roger Crisp and Corine Besson for providing some very helpful feedback on earlier drafts of this paper.

References

- Bengson, J., & Moffett, M. A. (2012). Two conceptions of mind and action: Knowing how and the philosophical theory of intelligence. In J. Bengson & M. A. Moffett (Eds.), *Knowing how: Essays on knowledge, mind and action*. Oxford: Oxford University Press.
- Gladwell, M. (2006). *Blink*. London: Penguin Books.
- Grimm, S. (2006). Is understanding a species of knowledge? *British Journal for the Philosophy of Science*, 57, 515–535.
- Hawthorne, J. (2003). *Knowledge and lotteries*. New York: Oxford University Press.
- Hills, A. (2009). Moral testimony and moral epistemology. *Ethics*, 120, 94–127.
- Hills, A. (2010). *The beloved self*. Oxford: Oxford University Press.
- Kvanvig, J. (2003). *The value of knowledge and the pursuit of understanding*. New York: Cambridge University Press.
- Poston, T. (2009). Know how to be Gettiered? *Philosophy and Phenomenological Research*, 79, 743–747.
- Pritchard, D. (2005). *Epistemic luck*. New York: Oxford University Press.
- Pritchard, D. (2009). Knowledge, Understanding and Epistemic Value. In A. O’Hear (Ed.), *Epistemology*. New York: Cambridge University Press.
- Pritchard, D. (2010). Safety-based epistemology: Whither now?. In Steup, M., Turri, J., & Sosa, E. (Eds.), *Contemporary debates in epistemology*. Oxford: Wiley Blackwell.
- Stanley, J. (2007). *Knowledge and practical interests*. Oxford: Oxford University Press.
- Stanley, J. (2011a). Knowing (How). *Nous*, 45, 207–238.
- Stanley, J. (2011b). *Know how*. Oxford: Oxford University Press.