

Moral Understanding as Knowing Right from Wrong¹

Moral understanding is a valuable epistemic and moral good. I argue that moral understanding is the ability to know right from wrong. I defend the account against challenges from non-reductionists, such as Alison Hills, who argue that moral understanding is distinct from moral knowledge. Moral understanding, she suggests, is constituted by a set of abilities: to give and follow moral explanations and to draw moral conclusions. I argue that Hills' account rests on too narrow a conception of moral understanding. Amongst other things, it cannot account for the importance of first-personal experience for achieving moral understanding.

1. Introduction

Moral understanding is a valuable epistemic good. It's something we look for when deciding whom to rely on for moral advice. It's a goal of moral education: something we hope to instill in our children. And it's an integral part of moral wisdom: a moral sage is someone who has profound moral understanding.

It's natural to think of moral understanding and moral knowledge as closely related epistemic phenomena. We say that moral understanding is a matter of knowing right from wrong. Typical sources of moral knowledge – thinking hard about a moral question, or having first-hand experiences – also seem to be sources of moral understanding.

This intuitive thought has lately come under sustained attack. In a recent paper, Alison Hills argues that moral understanding “is an important epistemic state with a different role to that traditionally filled by propositional knowledge.”² Hills' approach is in tune with a general trend in epistemology, where there has been a lot of interest in how understanding relates to knowledge. The assumption that understanding differs in several

¹ This paper has been presented at the University of Nottingham, the University of Zurich, and the University of Bristol. I am grateful to the audiences there for their comments and questions, and to Alexander Bird, Sophie Horowitz, Rae Langton, and Katia Vavova. I am particularly grateful to Christoph Baumberger, Tom Dougherty, Stephen Grimm, Yongming Han, Nathan Horward, Bernhard Salow, an anonymous referee for this journal and several editors for taking the time to read a draft of this paper and give generous and helpful feedback.

² Allison Hills, “Understanding Why,” *Noûs* (forthcoming), 2

crucial ways from knowledge is a common starting point.³ Let's call the claim that moral knowledge and moral understanding are distinct *non-reductionism*.

Non-reductionists point to several motivations for their approach. They note the received wisdom that it's easier to come to know something than it is to come to understand it. They point out that we often attribute knowledge without attributing understanding. And, in particular, they suggest that, unlike understanding, knowledge can be easily acquired and passed on through testimony. Thus, as Hills argues:

[A] centrally important concept in moral epistemology is not moral knowledge, but what I call *moral understanding*, and that the latter relates to testimony and to expertise quite differently from the former.⁴

What grounds this alleged difference between moral understanding and moral knowledge? Hills argues that, unlike knowing why p, understanding why p involves a range of abilities. When q is why p, understanding why p requires the abilities to:

- (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'.⁵

Hills argues that these abilities are both *necessary* and quite likely sufficient for having moral understanding:

These abilities are, I think, individually necessary for moral understanding, and I suspect that they may be jointly sufficient, provided that it is true that p and that q is why p (though I am open to the possibility that other abilities may be required in addition).⁶

³ For examples, see Jonathan Kvanvig, *The Value of Knowledge and the Pursuit of Understanding* (Cambridge: Cambridge University Press, 2002), Kate Elgin, "True enough," *Philosophical Issues* 14 (1)(2004), 113–131 and "Understanding and the facts," *Philosophical Studies* 132 (1) (2007), 33 - 42, Michael Strevens, "No Understanding without Explanation," *Studies in History and Philosophy of Science*, 44 (2013), 510–515, Linda Zagzebski, "Recovering Understanding," in *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*, ed. Matthias Steup (New York: Oxford University Press, 2001), Duncan Pritchard, "Knowing the answer, understanding and epistemic value," *Grazer Philosophische Studien* 77 (1) (2008), 325–339 and "Knowledge, Understanding and Epistemic Value," *Royal Institute of Philosophy Supplement* 84 (64) (2009), 19–43.

⁴ Alison Hills, "Moral Testimony and Moral Epistemology," *Ethics* 120 (1) (2009), 94–127, p. 97.

⁵ *ibid.*, 102–103

⁶ *ibid.*, 103

The abilities that Hills identifies all concern moral reasoning and reflection. On Hills' view then, understanding why an action is morally wrong requires the ability engage in moral reasoning about why it's wrong, where moral reasoning includes the ability to give and follow explanations. I will refer to this thought as the *moral reasoning claim*.⁷

My aim in this paper is to motivate and defend an alternative picture of moral understanding. The basic idea I shall develop is that moral understanding is the ability to know right from wrong. There are many ways in which we can achieve moral understanding: by perception, by first-personal experience, and even by moral testimony. In particular, agents can achieve moral understanding of why, for example, sexual harassment is morally wrong even when they lack the ability to articulate their understanding. This picture respects the intuitive thought that moral understanding and moral knowledge are closely related; it's reductionist. And, contra the moral reasoning claim, it offers a considerably broader conception of moral understanding.

Here's the plan: I will start by clarifying the terms of the debate. I suggest that we need to draw a distinction between specific instances (or achievements) of moral understanding and the capacity of moral understanding. This gives us a better grip on the central elements of Hills' account. I argue that non-reductionism is most plausibly taken as a claim about the instances of moral understanding while the moral reasoning claim targets the capacity of moral understanding. I then argue against these claims two separately. Section 3 and Section 4 make a case for reductionism. Section 5 argues that identifying the capacity of moral understanding with the capacity for moral reasoning is too narrow; in particular, it cannot account for the role of first-personal experience in moral understanding. These two strands of argument together motivate an alternative approach to moral understanding. I develop this account, the Moral Knowledge Account, in Section 6. Section 7 responds to a central objection to the Moral Knowledge Account.

2. Moral Understanding: Achievement versus Capacity

We need to distinguish between the capacity of moral understanding and particular instances of moral understanding. These two are importantly different and for the purposes of giving a philosophical account we need to

⁷ Hills is not the only proponent of the *moral reasoning claim*. Strevens, *No Understanding without Explanation*, 515 suggests that "to understand a moral rule might be to grasp a correct moral explanation of the rule (perhaps a derivation of the rule from fundamental moral principles) or to have the ability to use the rule to explain moral facts (for example, why lying is usually wrong)". More generally, philosophers of science have been emphasizing the relationship between understanding and explanation. Similarly, Kvanvig, *The Value of Knowledge*, 198 suggests that the ability to answer requests for explanation about a subject matter is constitutive of understanding. See, also, Kareem Khalifa, "The Role of Explanation in Understanding," *British Journal for the Philosophy of Science* 64 (1) (2013), 161-187. For a dissenting view, see Peter Lipton, "Understanding without explanation," in *Scientific Understanding: Philosophical Perspectives*, ed. Hans W. de Regt et al. (Pittsburgh: University of Pittsburgh Press, 2009), 43-63.

keep track of which one we have in mind. To credit someone with an instance of moral understanding is to say that they have achieved a particular *mental state* – for example, the mental state of understanding why one should not sleep with someone too drunk to give their consent or why one should not spank one’s children. In contrast, to credit someone with moral understanding *simpliciter* is to attribute to them a *capacity*: the capacity to achieve instances of moral understanding. This is just an instance of a more general distinction between an epistemic faculty and those states that are the output of this faculty. For example, we distinguish between color vision and seeing that the apple is red, or between sense of pitch and hearing that the note is c2. The former is a capacity or ability, the latter is a mental state acquired by exercising this capacity.

An agent can have the capacity of moral understanding even when they have not achieved some particular instance of moral understanding. Thus, for example, someone may fail to understand why the death penalty is wrong because they have been presented with misleading evidence about how it’s administered and its deterrent effect. They may nevertheless have the capacity of moral understanding: were they presented with non-misleading evidence, they would achieve moral understanding why the death penalty is wrong.

An agent can also have the capacity of moral understanding and fail to exercise it – either by choice or circumstances. But to have this capacity, she needs to be able to exercise it: there needs to be an existing cognitive mechanism that, employed correctly, will yield instances of moral understanding. Thus, a newborn does not (yet) have the capacity of moral understanding, just as the newborn does not (yet) have the capacity to speak English. Rather, she has the disposition to acquire and develop these capacities.

For clarity’s sake, I will use the terms “capacity of moral understanding” versus “instances of moral understanding” or “achievements of moral understanding” in what follows. (I’ll use the latter two interchangeably.) With this distinction on the table, let’s go back to non-reductionism and the moral reasoning claim.

In her discussion, Hills presents non-reductionism and the moral reasoning claim as a package deal: she argues that moral understanding involves abilities for moral reasoning by arguing that moral understanding differs from moral knowledge.⁸ But note that the two claims are independent. Consider first non-reductionism. Hills argues that moral understanding is an epistemic state distinct from moral knowledge. To defend it, she points to cases in which an agent purportedly understands why someone is evil without knowing why they are evil. Thus, it’s most natural to understand non-reductionism is as a claim about the nature of instances of moral understanding. Non-reductionism says that the mental state of understanding why an action is right and the mental state of

⁸ See Hills, *Moral Testimony and Moral Epistemology*. A similar argument is offered by Kvanvig, *The Value of Knowledge*, Strevens, *No Understanding Without Explanation*, and Pritchard, *Knowledge, Understanding, and Epistemic Value*.

knowing why it's right are two different mental states. An agent can be in the former without being in the latter.

The moral reasoning claim, on the other hand, is plausibly a claim about the capacity of moral understanding. It says which abilities are constitutive of this capacity: namely the abilities to give moral explanations, to draw the relevant conclusions, to justify one's conclusion in one's own words, etc. Thus, Eleanor's moral understanding of a particular moral subject matter – animal welfare, for example – is constituted by her ability to reason about this subject matter: to draw conclusions, justify them, give explanations.

To see that these claims are independent, notice that we can accept the moral reasoning claim and deny non-reductionism and *vice versa*. On the former, the capacity of moral understanding is the capacity for moral reasoning. An agent achieves an instance of moral understanding when she acquires moral knowledge by exercising her capacity for moral reasoning. The latter denies that the capacity of moral understanding is constituted by the ability for moral reasoning but nevertheless insists that instances of understanding are a distinct type of mental state.

The next two sections focus on instances of moral understanding, arguing in favor of reductionism. For this purpose, I will simply concede the moral reasoning claim is correct, i.e. that the capacity of moral understanding is the capacity to engage in moral reasoning.⁹

3. Defending Reductionism: Knowledge is Necessary for Achieving Understanding

We have seen that according to non-reductionism, achievements of moral understanding are different mental states from instances of moral knowledge. In contrast, reductionism is committed to the claim that when an agent understands why an action is right or wrong this is *in virtue of* having knowledge about why it's right or wrong. According to reductionism, knowledge is *all there is* to understanding; there is no need to stipulate a novel cognitive state that goes over and beyond knowledge. Reductionism thus makes a metaphysical claim. Just as physicalists about the mental say that every mental state is constituted by a physical state, so reductionists about understanding say that every instance of understanding is constituted by an instance of knowledge.

This section argues, contra Hills, that instances of knowledge are necessary for instances of understanding. The next section takes up the question whether knowledge is all there is to instances of understanding.

Hills' Argument from Epistemic Luck

Hills argues that knowledge is not necessary for achieving moral understanding, since instances of moral understanding have a different relationships to epistemic luck than knowledge. Hills appeals to the following case:

⁹ I revisit this concession in Section 5.

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.¹⁰

According to Hills, you don't know why Stalin was evil, since you do not know that he killed millions. Nevertheless, Hills suggests, you understand why Stalin was evil. And so, an agent can have moral understanding even in the absence of moral knowledge. More generally, Hills argues, following Kvanvig and Pritchard, that understanding is immune to environmental epistemic luck – the kind of luck exemplified in fake barn cases.¹¹

Let's grant that in this case you do not know why Stalin is evil.¹² But why should we accept that you nevertheless understand why he's evil? Hills argues that you understand why he is evil because you formed your true belief that Stalin was evil by exercising your capacity for moral reasoning. Moral understanding is the capacity to engage in moral reasoning. And if you form a true belief about why Stalin is evil by exercising the capacity of moral understanding, you thereby understand why Stalin is evil. She argues:

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*.¹³

But there are good reasons for rejecting the crucial premise: that if you form a true belief why *p* by exercising the capacity of moral understanding, you thereby understand why *p*. If this principle were true, we should expect it to generalize to other epistemic faculties. But it does not generalize. For example, forming a true belief that *p* based on exercising one's visual perception, does not guarantee that one sees that *p*. Suppose that Mary forms a belief that there is a red apple in front of her, based on her visual perception. But the object in front of her is an apple-shaped box, rather than a real apple. As it happens, however, there is, hidden from Mary's view, a red apple inside the box. Mary's belief is true. She arrived at it

¹⁰ Hills, *Moral Testimony and Moral Epistemology*, 204. Hills argument mirrors those made earlier by Kvanvig, *The Value of Knowledge*.

¹¹ The argument below also responds to those who argue that understanding does not require knowledge because it is not factive, e.g. Elgin, *True Enough*.

¹² This is not uncontroversial. See Stephen Grimm, "Is understanding a species of knowledge?" *British Journal for the Philosophy of Science* 57 (3) (2006), 515-535 and Paulina Sliwa, "Understanding and Knowing," *Proceedings of the Aristotelian Society* 115 (2015), 57-74.

¹³ Hills, *Moral Testimony and Moral Knowledge*.

by exercising her perceptual capacities. But it doesn't follow that she thereby *sees* a red apple or that she *sees that* there is a red apple in front of her.¹⁴

In the same way, we can say that the agent has exercised her capacity of moral understanding and that she has thereby acquired a true belief about Stalin, while denying that this true belief is an instance of moral understanding. The reductionist can grant that we are inclined to intuitively attribute moral understanding to the agent. But, she argues, this attribution concerns her *capacity* for moral understanding. After all, we have seen this capacity on display: she has arrived at the conclusion that Stalin is evil by correctly responding to the evidence at hand.¹⁵ Hills' epistemic luck case then does not give us reason to dismiss reductionism.

*Linguistic Data*¹⁶

There are positive reasons for accepting the reductionist claim that achievements of understanding require knowledge, based on linguistic considerations. These are relevant because non-reductionists do not purport to be stipulating a novel theoretical concept "understanding". Rather, they profess to pick up on a mental state that plays a pervasive and important role in our folk epistemic practice.

But if our folk epistemic practice recognizes a notion of understanding that's distinct from knowledge, we would expect this to be reflected in the semantics of "knowing" and "understanding." In particular, we would expect to be able to attribute one but not the other to an agent. But consider the following sentence:

I understand why Stalin was evil and I don't know why Stalin was evil.
This sounds infelicitous. Third-personal cases sound just as strange.
Consider:

Jane understands why she ought to give to charity and she does not know why she ought to give to charity.

The phenomenon generalizes to other forms of understanding. Understanding, like knowledge, takes *that*, *how*, and *wh*-clauses. Thus, consider:

¹⁴ If you are worried about deviant causal chains, a fake barn case will do as well. The agent obliviously driving through fake barn country and chancing upon a real barn may form a true belief that there is a barn in front of them, exercising their perceptual capacities. But it does not follow that they *see that* there is a barn there. Thanks to Stephen Grimm for pressing me on this point.

¹⁵ In Hills, *Understanding Why*, Hills gives an additional argument against knowing why p being necessary for understanding why p, by arguing that knowledge but not understanding is defeasible. The reductionist can respond using the same strategy as in cases of epistemic luck: what accounts for our intuitions that agents retain understanding in the face of epistemic defeat is that we, rightly, credit them with the capacity of understanding. Whether knowledge really is defeasible is controversial, see Maria Lasonen-Aarnio, "Unreasonable Knowledge," *Philosophical Perspectives* 24 (1) (2010), 1-21 and Higher-Order Evidence and the Limits of Defeat. *Philosophy and Phenomenological Research* 88 (2) (2014), 314-345.

¹⁶ My argument here builds on my previous work in Sliwa, *Understanding and Knowing*.

Steve understands what the right thing to do is and she does not know what the right thing to do is.

Hye-Sun understands how to do what's right and she does not know how to do what's right.

Fatma understands when stepping in is the right thing to do and she does not know when stepping in is the right thing to do.

Rahel understands that she ought to give to charity and she does not know that she ought to give the charity.

These all sound infelicitous, with the last one particularly jarring. The reductionist can readily explain this: since understanding *entails* knowing, these sentences are contradictory. They are semantically defective. The non-reductionist on the other hand owes us an explanation for the linguistic datum.

Could the non-reductionist simply argue that she does not share these linguistic intuition? This is not a promising response since evidence suggests that amongst competent native speakers they are widely shared. In a series of recent papers, Wilkenfeld, Plunkett, & Lombrozo argue that folk attributions of understanding systematically track attributions of knowledge. Summarizing their empirical studies on epistemic luck in particular, they report that there is “no evidence that people are more willing to attribute understanding than knowledge in cases involving epistemic luck”.¹⁷ The non-reductionist owes us an explanation of why this is so.

Perhaps the non-reductionist could grant that the sentences are odd but offer an alternative diagnosis. She might suggest that we should explain their oddity by appealing to pragmatic factors. Hills herself follows this route:

...that does sound odd, as does any sentence of the form: you understand why p but you do not know why p. Is this a serious objection to my interpretation of these cases and ultimately to my account of understanding? It would be if the sentence were completely incoherent. But I think that it is merely an unusual thing to hear.¹⁸

But the mere fact that a sentence is unusual cannot explain why it sounds odd. Most unusual sentences, while perhaps surprising, do not sound odd at all. (“The queen wears a hat made of pickles.”) A better explanation may go by way of conversational implicature. The non-reductionist could suggest that understanding does not semantically entail knowing but that to attribute an instance of understanding is to pragmatically imply that the other party knows. The mechanism here is presumably the same as in the case of sentences like “I ate some of the cookies”: this implies that there are some cookies left but it is, strictly speaking, compatible with my having emptied the cookie jar.

¹⁷ Daniel Wilkenfeld, Dillon Plunkett, and Tania Lombrozo, “Depth and deference: When and why we attribute understanding,” *Philosophical Studies*, 173(2)(2016), 373-393 and “Folk Attributions of Understanding: Is There a Role for Epistemic Luck?” *Episteme* (forthcoming).

¹⁸ Hills, *Understanding Why*, 17

Again it's hard to see how this will help the non-reductionist. Pragmatic implicature can be cancelled by explicitly denying what is standardly implied. Just as there is nothing odd about my saying "I ate some of the cookies. In fact, I ate all of them," we would expect there to be nothing odd about asserting "John understands why killing is wrong but he does not know why killing is wrong."

4. Is Knowledge all there is to Understanding?

Many non-reductionists accept that knowledge is necessary for instances of understanding.¹⁹ They contest the thought that knowing is *all there is* to understanding.

Before I get on to the arguments, one important clarification is in order. Non-reductionists often present their arguments as counterexamples to the claim that knowing why p is *sufficient* for understanding why p. They take reductionism to be characterized by the following principle:

An agent understands why p if and only if she knows why p.

But reductionists need not accept this principle.²⁰ What reductionists are committed to, first and foremost, is that every instance of understanding is constituted by knowledge. This is not to dismiss the non-reductionist arguments. They present the reductionist with three substantive challenges: to explain the relationship between instances of understanding and testimony, to explain how, on the reductionist picture, instances of understanding can come in degrees, and to explain why we often attribute to others knowledge without understanding.

The reductionism I defend here says that whether an agent has achieved understanding why p depends both on whether they know why p and on how much they know about whether p:

An agent understands why p *if and only if* she has a sufficient amount of knowledge why p.

I suggest that we should take amounts of knowledge to be a matter of which epistemic possibilities the agent recognizes. And I argue that how much knowledge is sufficient for an instance of understanding depends on context. I argue that the reductionist account of understanding meets each of the challenges. This puts it at an advantage *vis-a-vis* the non-reductionist, since the burden of proof generally lies on those who wish to make substantive additions to our ontology.

Hills' Argument from Moral Testimony

Hills argues that an agent who comes to know why an action is wrong by relying on testimony will generally lack understanding why the action is wrong:

¹⁹ For example, Stephen Grimm, "Understanding as knowledge of causes," in *Virtue Epistemology Naturalized: Bridges Between Virtue Epistemology and Philosophy of Science*, ed. Abrol Fairweather (London: Springer, 2014), 329-346.

²⁰ Though some do, for example, Amber Riaz, "Moral understanding and knowledge," *Philosophical Studies* 172 (1) (2015), 113-128 as well as an earlier time-slice of myself in Sliwa, *Knowing and Understanding*.

If you are attempting to gain knowledge, testimony can serve as the justification for your own belief but it is not usually a good way of acquiring moral understanding. Understanding why p will not – cannot – have the same relationship with testimony as knowing why p .²¹

Hills appeals to the case of Eleanor and Mary. Mary, who has thought about vegetarianism a lot and understands why it's morally required, tells Eleanor that it's wrong to eat meat because modern animal farming is cruel. Eleanor accepts what Mary says. Hills argues that while Eleanor has thereby gained testimonial knowledge why eating meat is wrong, there is nevertheless an epistemic asymmetry between Eleanor and Mary. The challenge for the reductionist is to explain this asymmetry.

Hills suggests that the best explanation for this persistent epistemic asymmetry is a difference in mental states: Mary understands why eating meat is wrong – she “grasps” the relevant proposition – while Eleanor merely knows why it's wrong – she “assents” to the proposition. She argues:

If this is correct [i.e. eating animals is indeed wrong because of the cruelty of modern farming] and Eleanor believes her, she knows not just that eating animals is wrong but she knows why too. But she is still not in the same position as Mary. Eleanor cannot draw relevant distinctions, cannot come to correct conclusions about similar cases. (What about animals reared under better conditions? What about fish?) Eleanor has been told why eating meat is wrong, but she does not really grasp the reasons why it is wrong.²²

Hills is right that, plausibly, Mary's moral testimony none withstanding, Eleanor and Mary are not epistemically on a par. But this asymmetry is something that the reductionist can both accommodate and explain.

The reductionist argues that we can account for the difference between Mary and Eleanor in terms of *how much* they know. Eleanor may well know that eating animals is wrong because of the cruelty of animal farming. But there is plausibly much that she does not know. For example, she may not know why it is cruel. Such knowledge requires some familiarity with how animals are kept and slaughtered. But Eleanor knows very little about these things. Mary, in contrast, has researched the issue extensively. And so, she plausibly knows not only that it's cruel but also why and how it's cruel.

According to the reductionist then, the difference between Eleanor and Mary is not a difference in the kind of cognitive attitudes that they bear to some fixed content – with Mary “grasping” and Eleanor “assenting” it. Rather Mary and Eleanor differ in what they know. They are in the same mental state – knowledge – with respect to different contents. What Eleanor needs to achieve epistemic parity with Mary is more knowledge.

²¹ Hills, *Moral Testimony and Moral Epistemology*, 121. A similar objection is widely made in the literature on understanding. See, for example, Pritchard, *Knowledge, Understanding and Epistemic Value*, Kvanvig, *The Value of Knowledge*, or Elgin, *True Enough*.

²² Hills, *Moral Testimony and Moral Epistemology*, 100.

But, the non-reductionists may object, if all that is needed is more knowledge, can't this be simply had by more testimony? Given enough expert testimony, you will eventually know what the expert knows. But, non-reductionists argue, there may still be an epistemic asymmetry between you and the expert. Thus, Grimm argues:

In reply, it might be said that this attitude of grasping [i.e. understanding] is nothing significantly different from an act of assent: it just involves more assent [i.e. more knowledge]. But it seems clear that one can pile up assents as high as you like without getting a grasping. [...] I can be told that the key to the proof lies in the fact that the first element is divisible by the second, and I can assent to this (again, based on the testimony of someone I trust). But I still might not understand the proof; I might fail to see or grasp how the truth of the theorem depends on the fact that the first element is divisible by the second. And it seems like this game could go on indefinitely. You might continue to spell out how the various dependencies are supposed to work. Moreover, based on your testimony I might assent to these claims at every step of the way. But none of these assents, in and of themselves, adds up to a genuine grasping on my part. A new kind of cognitive achievement seems to be needed.²³

This objection overestimates how easy it is to transmit knowledge by testimony. It relies on the thought that whenever the epistemic difference between two agents is a matter of what they know, this difference can be overcome by testimony. But this assumption faces multiple challenges.

First, to testify, you have to put your knowledge into words. But the fact that the expert knows that *p* does not guarantee that she will be able to articulate her knowledge: that she will be able to express it in a sentence whose meaning captures exactly what she knows. The worry here is not simply that people differ in how articulate they are. It's that what we can communicate is systematically constrained by our language. And at times the tools available in our language are blunt instruments when compared to the nuance of what we know. Thus, the art historian may know that a painter used a very particular shade of blue – she can picture it in her mind, she immediately recognizes it when she sees it – but our language may lack the precise term for it. The best she can do is to describe (“it’s very intense and brilliant, with just a touch of green, but not too much...”). But no matter how lengthy, the description will be unlikely to exactly nail down the shade in question. This difficulty arises in the case of moral knowledge, too. Think about a victim of stalking before the term entered common usage in the early 1990s. The victim may herself know quite well that and how she is

²³ Stephen Grimm, “Is understanding a species of knowledge?” *British Journal for the Philosophy of Science* 57 (3) (2006), 515-535, particularly p. 532. Thank you to an anonymous reviewer for drawing my attention to this objection.

being wronged; nevertheless, she may not be able to fully convey the nature of the wrong to others.²⁴

Even setting these difficulties aside, it's not true that by "assenting" to what the expert tells you, you always come to know what the expert knows. Assent is cheap; I can assent to what I have been told by simply believing that you spoke truly. But I can believe you spoke truly – that the sentences you have uttered express a truth – even if I don't know *what* you said: I don't know *which* truth you are expressing. What I can come to know based on your testimony, depends on many factors. It depends, in part, on my linguistic and conceptual resources. But it also depends on how much I already know about the subject matter at hand.

To see this, we need a better grip on what it is to have more or less knowledge than someone else. A natural first answer is this: knowing more is a matter of knowing *a greater number* of propositions. Elgin argues:

A subject's knowledge consists of discrete grains, each separately secured. She amasses more knowledge by accumulating more grains.²⁵ Those "grains" are "individual facts, expressed in true propositions and/or stated in true declarative sentences."²⁶ They can be individually transmitted in testimony. If you accept what I say, you acquire the "grain." And so it goes, "grain by grain" until eventually we know all the same things. We see this conception at work in Grimm's quote above: coming to know more about the mathematical proof just is coming to assent to a *greater number of claims*.

But this is not a plausible conception of what different amounts of knowledge consist in. To get a sense of the difficulties, consider how many "grains" someone acquires when they learn that it was a dark and stormy night. One? Or three? ("it was dark", "it was stormy", "it was night"). But they also learn that it was not a light and sunny day. Do these count separately? Trying to determine how much someone knows by counting sentences or "bits" of content is a non-starter. Content does not come pre-carved into individual discrete bits.²⁷

²⁴ The phenomenon I have in mind here has important similarities to what Miranda Fricker calls hermeneutical injustice. See her *Epistemic injustice: Power and the ethics of knowing* (New York: Oxford, 2007).

²⁵ Catherine Elgin, "From Knowledge to Understanding," in *Epistemology Futures* ed. Stephen Cade Hetherington (New York: Oxford University Press, 2006), 199-215, p. 203. See also Georgi Gardiner, "Understanding, Integration, and Epistemic Value," *Acta Analytica* 27 (2) (2012), 163-181. It's worth pointing out the assumption has proponents outside non-reductionism.

²⁶ Elgin, *From Knowledge to Understanding*, 203.

²⁷ This line of argument is developed in much more detail in Nick Treanor, "The Measure of Knowledge," *Noûs* 47 (3) (2013), 577-601 and "Trivial Truths and the Aim of Inquiry," *Philosophy and Phenomenological Research* 89 (3) (2014), 552-559. The point has also been raised by Robert Stalnaker, *Inquiry* (Cambridge: MIT Press, 1987), 64-65, Daniel Dennett, *Content and Consciousness* (New York: Routledge, 1969), 183, Michael Williams, *Problems of Knowledge* (Oxford: Oxford University Press, 2001), 131 and David Lewis, "Reduction of mind," in *Papers in Metaphysics and Epistemology, Volume Two*. (Cambridge: Cambridge University Press, 1999), 311, amongst many others.

A more promising approach starts from the observation that knowing that things are thus-and-so involves knowing that they aren't otherwise. To know that p is to discriminate among various ways the world might be and to correctly locate the actual world on the right side of the divide.²⁸ I suggest that how much an agent knows is a matter of which epistemic possibilities they recognize and exclude.

To illustrate, recall Mary and Eleanor. Both know that eating animals is wrong because modern animal farming is cruel. But the epistemic possibilities among which Mary discriminates are much more fine-grained than those of Eleanor. Knowing what animal farming looks like and what animals need, she realizes that it can be cruel to animals in different ways: their stalls may be crowded and filthy, their natural movement constrained, their diet may be inappropriate, they may be separated from their young. Animals may suffer on the way to slaughter, they may be terrified, or killed in a painful way. Animal farming may also be cruel to the workers in the industry. It may pollute the environment and affect wildlife. Mary recognizes these as different ways in which animal farming could be cruel and she knows which ones obtain.

Eleanor has some idea of what cruelty involves, she recognizes that the animals may be in pain or frightened or hungry. Based on Mary's testimony, she knows that one of these obtains, but not which one. But she cannot conceive of more specific possibilities. And insofar as she has some misconceptions of what animal farming looks like, she may also entertain some possibilities that only she but not Mary entertains.²⁹

Eleanor and Mary both know that animal farming is cruel not because they have some "grain" in common but rather because there is overlap between the epistemic possibilities they recognize for what animal farming could be like. But they differ in what exactly they know: Eleanor's range is more limited, less fine-grained, in some ways and perhaps more inclusive in others.

This is, of course, just a sketch. Developing these remarks into a full account of how to quantify knowledge would lead us too far away from the topic at hand.³⁰ But this outline is enough to make clear why we should not expect testimony to be a panacea for epistemic asymmetries. What a hearer can learn from a testimonial exchange depends as much on the hearer as on the speaker; it depends on which epistemic possibilities the hearer recognizes.

²⁸ This observation is found, amongst others, in Stalnaker, *Inquiry* and David Lewis, "Elusive knowledge," *Australasian journal of Philosophy*, 74(4) (1996), 549-567.

²⁹ More knowledge is thus not simply a matter of recognizing more epistemic possibilities. Sometimes it can be a matter of recognizing less. You distinguish between the possibilities that either Clark Kent or Superman saved the child. I am in on Superman's secret, so I don't.

³⁰ The central idea here is similar to Treanor's account in Treanor, *Measure of Knowledge*. Treanor develops this proposal in more detail, suggesting that we can use a similarity relation to quantify knowledge.

To illustrate, suppose that Jones (who has no medical expertise) and Smith (a medical doctor) are talking about a recently deceased friend and wonder why he died. A mutual acquaintance who, like Smith, is a medical doctor informs them that it was lung cancer: “Small cell carcinoma. Nasty stuff.” Jones and Smith get the same testimony. Both accept the doctor’s words; they believe him.

Yet, they plausibly differ in what they have come to know because they differ in which epistemic possibilities the testimony allows them to exclude. Jones comes to exclude the possibilities that his friend died from a car crash, a heart attack, and cancers of various other body parts. Smith eliminates some of the same possibilities – the car crash, the brain tumor, the heart attack – as well as that it was large cell lung cancer and mixed cell carcinoma. Jones cannot entertain these because she lacks both the relevant conceptual resources and medical background knowledge. While both Jones and Smith believe what the speaker tells them and accept her words, they nevertheless differ in what they come to know on the basis of the testimony.

This should not be surprising. Giving and taking testimony falls under the general category of giving and accepting assertions. And it’s familiar that the effects of any particular assertion depend on what *both* speaker *and* hearer bring to the table. Thus, as Stalnaker argues:

...the essential effect of an assertion is to change the presuppositions of the participants in the conversation by adding the content of what is asserted to what is presupposed.³¹

The presuppositions are the epistemic possibilities that the participants of the conversation entertain – the ways the world might be, according to those agents. When there is a significant asymmetry in expertise, the presuppositions of speaker and hearer will generally not align: the expert can discriminate among ways the world might be that the lay person is oblivious to and the lay person may discriminate among possibilities that the expert recognizes as one and the same. And since, for example, one cannot eliminate a possibility that one cannot entertain, the hearer will generally not come to have the same knowledge as the speaker. Insofar as she accepts the speaker’s words, the hearer will learn something. But she will learn only part (and perhaps only a small part) of what the speaker said. In the extreme case, the hearer may only learn that the sentence uttered by the expert expresses a truth, while remaining entirely ignorant of which truth it expresses.

The reductionist thus agrees that, when there is a significant epistemic asymmetry between two parties, transmitting instances of understanding via testimony may not be easy. But she allows that in some cases it is possible: by relying on enough testimony, we can sometimes come to know what the expert knows, or at least, we can come to know enough to achieve understanding. In some cases, it may take a lot of testimony – too

³¹ Robert Stalnaker, “Assertion,” in *Context and content: Essays on intentionality in speech and thought*. (Oxford: Oxford University Press, 1999), 85. We don’t have to rely on Stalnaker’s framework here, the same point can be made drawing on David Lewis, “Scorekeeping in a language game,” *Journal of Philosophical Logic* 8 (1)(1979), 339-359.

much for a single exchange. Thus, imagine Eleanor and Mary having many, many conversations about animal farming over the course of several weeks. Eventually, Eleanor could learn enough from Mary to achieve understanding of why eating meat is wrong. And by going to medical school and listening to a few years worth of lectures in anything from biochemistry to oncology, Jones' medical knowledge might match that of Smith. In others, the epistemic asymmetry might be unsurmountable by any finite amount of testimony.

Let me say more in defense of this last thought: that achieving an instance of understanding through testimony is possible. For you might suggest an alternative view: instances of moral understanding are instances of moral knowledge *based on moral inference*. What is it for a knowledge to be based on moral inference? A natural idea is that the agent must have herself inferred that, say, the action is right from its right-making features.³² For example, she must have inferred the moral status of eating meat from facts about how animals are kept. On this view, it's not the quantity of knowledge that matters but its quality. The knowledge needs to be of the right kind, it needs to be based on a moral inference.³³

But drawing the relevant inference yourself is neither necessary nor sufficient for achieving an instance of understanding. To see that it's not sufficient, consider Yasmin, who is joining a group of animal rights activists. Members of this group believe that eating meat is wrong based on very extensive research into modern animal farming, slaughter and, animal wellbeing. Yasmin, too, believes that eating meat is wrong. But her belief is based on a fragment of a documentary she once listened to while driving. She's hazy on the details but she remembers it mentioning that modern animal farming was cruel to animals. From this, she inferred that it's morally wrong to eat animals. Since it's the cruelty of animal farming that makes eating meat wrong (we can suppose), Yasmin's belief is based on a moral inference. But, given her limited knowledge when compared to the animal rights activists, we would not credit her with an instance of moral understanding.

More importantly, however, you do not need to draw the relevant inference yourself to achieve moral understanding. Imagine Samir, who faces a tricky situation. He just met his best friend's new boyfriend. His friend is completely infatuated with this guy. But, after seeing them together, Samir cannot shake a bad feeling. He is uncertain how to answer his friend's inevitable question about his opinion. Samir recognizes the various considerations that pull in different directions: his concern for his friend's wellbeing and their close friendship based on straight talk on the one hand, the risk of hurting his friend's feelings and endangering their friendship on

³² The right-making features constitute what Enoch calls "transparent moral evidence". See David Enoch, "A Defense of Moral Deference," *Journal of Philosophy* 111 (5) (2014), 229-258 at 216. Note that I'm not taking any stand here on what this inference involves: it could involve an episode of conscious reasoning but it could also be a sub-personal affair. I say a bit more about different accounts of moral reasoning in footnote 45.

³³ I'm grateful to an anonymous referee for suggesting this alternative.

the other. Samir sees the significance of each consideration but he is uncertain how to weigh them up. He turns to someone whose moral judgment he trusts who, rightly, tells him that the reasons to tell the truth (John's wellbeing, in particular) are weightier than those to remain silent. It seems that, with this help, Samir comes both to understand that he should tell the truth and why he should do so. But Samir's moral knowledge is not based on his own moral inference: it was precisely the moral inference that he needed help with.

Cases like Yasmin and Samir suggest that when it comes to the relationship between moral knowledge and instances of understanding, it's quantity that matters more than quality. Yasmin doesn't understand why eating meat is wrong because, even though she drew the inference herself, she does not know enough about why it's wrong. Samir does understand why he should tell the truth because, even though he didn't draw the inference himself, he knows a lot about why he ought to tell the truth.

The observation that knowledge comes in different amounts helps reductionism in other ways. First, it helps to address a longstanding challenge for the reductionist: how to make sense of the fact that achievements of understandings seem gradable when attributions of knowledge are binary. The answer is that the degree to which an agent understands why, for example, eating meat is morally wrong corresponds to how much she knows about why it's wrong. The more you know about why eating meat is wrong, the better you understand why it's wrong.

Second, it helps the reductionist give a satisfying explanation for the relationship between degrees of understanding and abilities. As Hills notes, the epistemic difference between Eleanor and Mary will manifest itself in what it is that they can do: which questions they can answer, what they can explain, etc. Unlike Mary,...

...Eleanor cannot draw relevant distinctions, cannot come to correct conclusions about similar cases. (What about animals reared under better conditions? What about fish?)³⁴

Hills suggests that, rather than taking this difference in abilities as something to be explained, we should rather take them as constitutive of instances of understanding: what it is for an agent to lack these abilities just is for them to lack an instance of understanding.

In contrast, the reductionist can explain why agents in Eleanor and Mary's position differ in what they can do: they differ in what they can do because they differ in what they know. Given that, compared to Mary, there is much about animal farming that Eleanor does not know, it's not surprising that there are many questions she cannot answer and many inferences she cannot draw. Eleanor lacks these abilities precisely because she lacks the requisite knowledge.

Hills' Linguistic Argument

Hills also appeals to linguistic considerations to argue against reductionism. She argues that non-reductionism offers us a straightforward

³⁴ Hills, *Moral Testimony and Moral Epistemology*, 100.

explanation of two distinct phenomena. The first is that English is not unique in having different words for “to understand” and “to know”. For example, German distinguishes between “wissen” and “verstehen”, French between “savoir” and “comprendre”, Finnish between “tietää” and “ymmärtää”, and Polish between “wiedzieć” and “rozumieć”.³⁵ The second is that there is nothing odd about attributing an instance of knowledge to someone without attributing understanding to them. Thus, consider the following sentence:

Eleanor knows why eating meat is wrong but she doesn't understand why eating meat is wrong.

This seems like a perfectly fine thing to say. Hills argues that the non-reductionist has a ready explanation: it's not surprising that languages use different words to attribute instances of knowledge and instances of understanding...

...because understanding why p and knowledge why p are separate states that play different epistemic roles.³⁶

And so, it's also not surprising that it can be felicitous to attribute knowledge without attributing understanding.

But the reductionist has a sound response to this challenge. The reductionist says that instances of moral understanding are constituted by moral knowledge. But not just any amount will do: you need to know a sufficient amount. What counts as sufficient depends on the salient alternatives in that context: this may be another agent and how much they know. Or it may be you at a different point in time – how much you knew or will know *then*.³⁷

This yields a contextualist semantics of understanding: the truth conditions of a sentence ascribing understanding to an agent vary depending on the context in which the sentence is used. “Understanding” is, in this respect, similar to the term “empty”. We can truly say “this drawer is empty” when the salient alternative is one that's full of junk. But the same sentence is false when uttered in a context where the salient alternative is a vacuum chamber. The context determines how little stuff there must be in order for a container to count as empty. Similarly, the context determines just how much knowledge an agent needs to have in order to count as having an instance of understanding.³⁸

Such contextualism straightforwardly accommodates the linguistic data that Hills points to. Consider:

³⁵ Though this may not be universal. See Riaz, *Moral Understanding and Knowledge*.

³⁶ Hills, *Understanding Why*, 9.

³⁷ A contextualist account of the semantics of understanding has also been suggested by Christoph Kelp, “Understanding phenomena,” *Synthese*, 192(12) (2015), 3799-3816. Kelp's concern however is to give an account of understanding of phenomena (what I refer to as the capacity of understanding), as opposed to individual instances of understanding.

³⁸ Note that this proposal is neutral on whether knowledge is itself a context-sensitive term, as for example, argued by Lewis, *Elusive Knowledge*.

Eleanor knows why she should not eat meat but does not understand why she should not eat meat.

Can we explain why this sentence seems both assertable and true? We can: it is true because the salient alternative here is Mary who is much more knowledgeable than Eleanor. Mary's epistemic situation sets the standard for how much knowledge is required in order to count as having an instance of understanding.

We can also explain why we use distinct words for "knowledge" and "understanding". The concept of understanding fulfills a socially useful function because it's useful to track how much knowledge individuals have on particular questions. Such information is particularly valuable when deciding who you should turn to for testimony and advice. If Mary knows a lot more about why eating animals is wrong than Eleanor, then she is very likely to know not just whether it's cruel but also what makes it cruel. Someone who has achieved understanding on a question will in general be a better source of knowledge because there is more knowledge that can be gained from her.³⁹

This kind of contextualism is also independently attractive. First, it explains why our attributions of understanding are sensitive to how much salient expertise there is in the background. Thus, it's fine for me to claim that I understand why smoking causes lung cancer at a dinner party with my philosophy friends. But it seems less fine when I'm a lay person at a convention of pulmonologists.

Second, contextualism explains why the following seem like good inferences:

If Eleanor knows why eating meat is wrong but doesn't understand why it's wrong, then there is something that Eleanor does not know about why it's wrong to eat meat.

If Mary understands why eating meat is wrong but Eleanor does not, then there is something that Mary knows about eating meat that Eleanor does not know.

These are natural inferences to draw since to say that Mary, but not Eleanor, has achieved understanding with respect to a question just is to say that Mary knows more than Eleanor about it. But they present a puzzle for non-reductionists, since agents can differ in what they understand even when there is no difference in what they know. But then, we should not be able to infer from a lack of understanding to a lack of knowing.

Third, contextualism explains why, once you have attributed a very large amount of knowledge to someone on a particular question, there is no

³⁹ Similarly, Wilkenfeld, Plunkett & Lombrozo, *Depth and deference*, 374, argue that the difference in attributing instances of knowledge versus instances of understanding "is related to speakers' interest in tracking who would be a valuable expert to consult about a particular field of inquiry." A principled way to develop this account would take Edmund Craig's discussion of the function of the concept of knowledge as its starting point. See Edmund Craig, *Knowledge and the State of Nature: An Essay in Conceptual Synthesis* (Oxford: Oxford University Press, 1990) as its starting point. I'm grateful to an anonymous referee for raising this challenge.

need to also attribute understanding to them. Thus, while not contradictory, the following sounds decidedly odd:

Lucy has an impressive amount of knowledge about the extinction of the dinosaurs – she knows everything there is to know, way more than any other scholar. And she also understands why dinosaurs became extinct.

This is unsurprising: to attribute an instance of understanding why dinosaurs became extinct *just is* to attribute a large amount of knowledge about this question to someone. Once you have done the latter, the outright attribution of understanding becomes superfluous; it violates the Gricean maxim of relevance. But on non-reductionism it should be surprising: since instances of knowing and instances of understanding are distinct mental states that can and (as we have seen) often do come apart, there is always the possibility of having the former without the latter.

Let's step back for a moment and take stock. The last sections made a case for taking knowledge to be necessary for instances of understanding. This section responded to non-reductionist arguments that knowledge cannot be all there is to instances of understanding. I have argued for a reductionism account on which every instance of understanding is constituted by knowledge. But whether an agent, for example, understands why *p* depends not just on whether they know why *p* but also on how much they know. I argued that such a reductionist has compelling explanations of the allegedly problematic phenomena. These explanations are not *ad hoc*; their central elements – what constitutes different amounts of knowledge, how testimony transmits knowledge, a contextualist account of the semantics of understanding attributions – are all independently plausible and well motivated. Reductionism thus gives us an account of instances of understanding that accommodates and explains its central features and is, at the same time, parsimonious. Thus, we should be reductionist about instances of moral understanding even if we concede the moral reasoning claim – that the capacity of moral understanding is the capacity for moral reasoning. The next section turns to the capacity of moral understanding and takes aim at this concession.

5. Moral Understanding without Moral Reasoning

According to the moral reasoning claim, the capacity of moral understanding is characterized by the abilities to follow moral explanations, draw the right moral conclusions based on your moral evidence, and give moral explanations yourself. This section argues that we should not accept the moral reasoning claim; it yields an implausibly narrow conception of the capacity of moral understanding. I will bring this out in two ways. First, I argue that the moral reasoning claim conflates having moral understanding and having the ability to articulate it. But these are distinct abilities and they can come apart. Second, I argue that the moral reasoning claim cannot accommodate the significance of first-personal experience for moral understanding.

Consider Julie, who is, by all accounts, a morally good person: honest, caring and kind. She is also extremely morally competent; for any

moral situation she finds herself in, she does the right thing because it strikes her as the right thing to do, and it strikes her as the right thing to do because she responds to the morally relevant features of the situation. I think Julie would rightly strike us as someone who has deep moral understanding. But Julie need not, also, have the abilities to give moral explanations, or provide an informative justification of her verdicts. She may not be very good at drawing conclusions about abstract moral scenarios. When pressed for an explanation, she might just shrug her shoulders and offer: “it’s just the right thing to do.” Or perhaps: “it would be unkind not to help.” But she knows right and wrong when she sees it. Julie’s predicament is just a more extreme version of a familiar situation; we often reach moral verdicts without any reflection and cannot give much of an explanation for why they are right. (Why is it wrong to drown babies in Coca-Cola for fun? Why is it wrong to harm someone? Well, *it just is.*)⁴⁰

I have suggested that agents like Julie have the capacity of moral understanding even when they do not have the kinds of abilities that the moral reasoning claim regards as essential for it. But you might insist that, while Julie does strike us as morally admirable in *some way*, we should nevertheless stop short of crediting her with the capacity for moral understanding. The problem is, you might argue, that the capacity of moral understanding requires more of an agent than just being in a position to do the right thing herself. Part of what it is to have this capacity is to be a good moral advisor. But one can only give moral advice if one can articulate and explain one’s moral insight. To be a good moral advisor, Julie thus needs to have the abilities outlined by the moral reasoning claim.⁴¹

There is something right about the general line of thought: someone who has great moral understanding (the capacity) is typically someone we can rely on for moral guidance. Moral understanding is not just personally but interpersonally valuable. But Julie can be an excellent source of moral guidance even when she lacks the abilities outlined by the moral reasoning claim. First, moral guidance need not take the form of verbal moral advice. We can learn not just from what others tell us but also from what they *show* us. Even if Julie is no good at articulating moral principles and explanations, she can show others what the right thing to do is, she can show them how

⁴⁰See Peter Railton, “The affective dog and its rational tale,” *Ethics* 124 (4) (2014), 813-859. The suggestion that our capacity for understanding can outstrip our capacity to articulate has been made in the context of other debates, particularly in discussions of particularism. See, for example, Paul Churchland, “The Neural Representation of the Social World,” in *Mind and morals: Essays on ethics and cognitive science*, ed. Andy Clark et al. (Cambridge: MIT Press, 1996), Joseph Raz, “The Truth in Particularism,” in *Engaging Reason* (Oxford: Oxford University Press, 1999), 218-246, Pekka Väyrynen, “Usable Moral Principles,” in *Challenging Moral Particularism*, ed. Mark Norris Lance (New York: Routledge, 2008), 75-106, and Larry Wright, “Reasons and the Deductive Ideal,” *Midwest Studies in Philosophy* 23 (1999), 97-206. I’m grateful to an editor of this journal for pointing me to this literature.

⁴¹ This argument is suggested by a discussion in Alison Hills, “The Intellectuals and the Virtues,” *Ethics* 126 (1) (2015), 7-36, at 28-29, where she argues that virtue requires the ability to engage in moral reasoning.

to be kind, caring, honest, how to resolve conflicting moral demands. In other words, she can offer her friends moral guidance by being a moral exemplar – a role model. Her friends can come to see what the right thing to do is in a situation by imagining how Julie would handle it. Moral guidance is not the same thing as verbal instruction.

Second, suppose you turn to Julie for help with how to deal with a difficult and quarrelsome relative at an upcoming family reunion. Should you make a scene? Or should you let things go? Julie does not embark on a lecture. Rather, she listens patiently to your story. Then she asks a number of insightful questions about what you think you should do, about your “problem relative”, about how conflicts typically unfold, about who else will be present. She doesn’t offer much in terms of moral explanations – she doesn’t even tell you outright what to do – but her questions draw your attention to certain morally relevant features of the situation that were not on your radar. You leave the conversation with real moral insight: you know what you should do. You come to see that there are other options to diffuse the provocation without causing other family members needless distress.

Julie has provided you with something very much like moral advice based on her capacity for moral understanding. She resolved your uncertainty about both what the right thing to do is and why it’s the right thing to do. But doing so did not require Julie to articulate any moral explanations, or to draw moral conclusions about related cases. The abilities involved in the moral reasoning claim are hence not necessary for moral guidance.

Such cases bring out that there is a difference between having the capacity of moral understanding and, for example, the ability to articulate moral explanations. An agent can have the former without having the latter.

⁴² Of course, someone’s being very good at giving moral explanations and reasoning through cases, can be excellent evidence for their capacity of moral understanding. But it does not constitute this capacity.⁴³

There is a second way in which the moral reasoning claim falls short as an account of the capacity of moral understanding: it cannot account for the significance of first-personal experience. Having a certain experience – witnessing a morally significant situation first-hand – can expand one’s capacity of moral understanding. It can also be the basis for achieving instances of moral understanding. Thus, consider George Orwell’s account of witnessing an execution in Burma first-hand:

⁴² Conversely, one can have the abilities involved in the moral reasoning claim while having a very limited capacity of moral understanding. One does not have to look far for real life examples: philosophers, excellent at reasoning about even complicated moral cases involving fairness and equality, and yet blind when it comes to instances of e.g. gender inequality right under their nose.

⁴³ Stephen Grimm also puts pressure on an articulacy requirement on understanding, pointing to cases of young children. See his “Understanding and Transparency,” in *Explaining Understanding: New Essays in Epistemology and the Philosophy of Science*, ed. Stephen Grimm et al. (New York: Routledge, 2017).

It was about forty yards to the gallows. I watched the bare brown back of the prisoner marching in front of me. He walked clumsily with his bound arms, but quite steadily, with that bobbing gait of the Indian who never straightens his knees. At each step his muscles slid neatly into place, the lock of hair on his scalp danced up and down, his feet printed themselves on the wet gravel. And once, in spite of the men who gripped him by each shoulder, he stepped slightly aside to avoid a puddle on the path.

It is curious, but till that moment I had never realized what it means to destroy a healthy, conscious man. When I saw the prisoner step aside to avoid the puddle, I saw the mystery, the unspeakable wrongness, of cutting a life short when it is in full tide. This man was not dying, he was alive just as we were alive. All the organs of his body were working - bowels digesting food, skin renewing itself, nails growing, tissues forming - all toiling away in solemn foolery. His nails would still be growing when he stood on the drop, when he was falling through the air with a tenth of a second to live. His eyes saw the yellow gravel and the grey walls, and his brain still remembered, foresaw, reasoned - reasoned even about puddles. He and we were a party of men walking together, seeing, hearing, feeling, understanding the same world; and in two minutes, with a sudden snap, one of us would be gone - one mind less, one world less.⁴⁴

It is natural to say that witnessing the execution led Orwell to understand that capital punishment is morally wrong and why it's wrong - or to understand it better. It is also natural to say that it expanded his capacity of moral understanding. But the moral reasoning claim conflicts with both of these things.

It conflicts with the former because it implies that agents can achieve instances of moral understanding only by engaging in moral reasoning (i.e. by drawing moral conclusions based on the relevant information, by articulating moral explanations, etc). This is because we achieve instances of moral understanding by exercising our capacity of moral understanding. According to the moral reasoning claim, this capacity just is the capacity to engage in moral reasoning. But Orwell's moral insight is not based on moral deliberation. He didn't reason his way to the conclusion that the death

⁴⁴ George Orwell, "A Hanging," in *George Orwell Essays* (London: Penguin Modern Classics, 2000).

penalty is morally wrong.⁴⁵ Rather, it is based on something more like a perceptual experience. He *saw* “the unspeakable wrongness” of killing another human being.

Orwell’s experience may also have expanded his capacity of moral understanding. After all, certain experiences can improve one’s epistemic position vis-à-vis a particular moral subject matter. But plausibly insofar as they broaden one’s capacity of moral understanding, it’s not because they make one better at moral reasoning. A disabled person may have a better understanding of ableism than an able-bodied person because of her experience of being disabled. But her deeper moral understanding need not be a matter of giving moral explanations. Rather, it may be a matter of being able to recognize ableism when one sees it, of giving certain considerations the moral weight they deserve, of noticing particular injustices. Similarly, working in a particularly hostile environment may lead a woman to have a better understanding of sexual harassment. But her increased capacity of moral understanding need not involve being better at drawing conclusions about cases. Rather, she may simply be in a better position to recognize sexual harassment both when she is on the receiving end of it and when others are, to notice little inequities that may have otherwise escaped her attention and giving them the correct moral weight. Her experience makes things salient to her that may not have been salient to her before.

The two lines of objection both point to one underlying problem with the moral reasoning claim: it offers a much too narrow conception of the capacity of moral understanding. Agents can expand and deepen their capacity of moral understanding with respect to specific moral questions even when they do not improve their abilities to “follow explanations why p” or “give explanations why p in their own words”. And they can achieve instances of moral understanding without exercising their capacity for moral reasoning. This is not to say that the ability to engage in moral reasoning is irrelevant to the capacity of moral understanding – it’s just to say that the capacity of moral understanding includes much more than that.

⁴⁵ But didn’t this involve something like moral reasoning? Some philosophers (for example, Ralph Wedgwood, “The normative force of reasoning,” *Noûs*, 40 (4) (2006), 660-686, Nomy Arpaly, *Merit, meaning, and human bondage: An essay on free will* (Princeton: Princeton University Press, 2009), Nomy Arpaly and Timothy Schroeder, “Deliberation and acting for reasons,” *Philosophical Review*, 121(2) (2012), 209-239) argue that reasoning can be merely a matter of causal transitions between mental states, which need not be mediated by conscious or deliberate mental actions; it’s not so much something we do as something we undergo. I don’t deny that Orwell underwent moral reasoning *in this sense* – after all, he formed a belief that killing the man was wrong based on his perception that it was wrong. The account of the capacity of moral understanding defended in the next section – the Moral Knowledge Account – is compatible with the thought that moral understanding involves reasoning in this minimal sense. But those sympathetic to the moral reasoning claim (for example, Alison Hills, Michael Strevens, Jonathan Kvanvig) explicitly take the capacity of understanding to involve abilities that constitute reasoning in a much more robust sense, on the personal, rather than the sub-personal level. My argument here targets these views. Thanks to Stephen Grimm and an anonymous referees for asking me to clarify this point.

Let's pause for a moment and review the argument so far. I started out with Alison Hills' account of moral understanding. I have argued that her account is characterized by commitments: non-reductionism, a view about the nature of those mental states that comprise instances of moral understanding, and Moral Reasoning, an account of the capacity of moral understanding. I have argued that these commitments are distinct and independent and that there are persuasive reasons to reject them both. These negative conclusions prepare the ground for the positive project of the paper. The next section develops an account of the capacity of moral understanding that incorporates the two lessons drawn out so far: the lesson that instances of moral understanding reduce to instances of moral knowledge, and that there is more to the capacity of moral understanding than a set of abilities related to moral reflection.

6. Moral Understanding as Knowing Right from Wrong

I suggest that the capacity of moral understanding is the *ability to acquire moral knowledge*. The ability to acquire moral knowledge is *constitutive* of moral understanding. So, an agent has moral understanding if and only if (and to the degree to which) she has the ability to acquire moral knowledge. I will call this the Moral Knowledge Account.

Let me start by unpacking its central elements. Following Williamson and others, I assume that knowledge is a mental state.⁴⁶ It is a personal, rather than a sub-personal state. This means that the agent must have some first personal access to its content; it must be available to guide the agent's actions. Moral knowledge is knowledge with a particular content; it employs moral concepts. Examples of moral knowledge include knowing *that* an action is right or wrong, just, fair, sexist, racist, or kind, knowing *what* the right (wrong, just, fair, sexist,...) thing to do is, knowing *why* an action is right (or wrong, just, fair, sexist,...), etc.⁴⁷

To have the ability to acquire moral knowledge, the agent needs to have a psychological mechanism (or plausibly, a set of psychological mechanisms) which, if it works correctly, yields moral knowledge when the agent is in the right circumstances – in particular, when she is presented with the relevant moral evidence.⁴⁸ An ability to acquire moral knowledge

⁴⁶ See Timothy Williamson, *Knowledge and its Limits* (Oxford: Oxford University Press, 2000), particularly chapter 1. For arguments in support of this claim see also John Gibbons, "Knowledge in action," *Philosophy and Phenomenological Research* 62 (3) (2001), 579-600 and Jennifer Nagel, "Knowledge as a Mental State" *Oxford Studies in Epistemology* 4 (2013), 275-310.

⁴⁷ While I'm sympathetic to intellectualist views of knowledge how and knowledge-wh, the Moral Knowledge Account is neutral on the question.

⁴⁸ Recall that on the notion of ability at issue here, I do not have an ability to, say, speak French merely because I could, in principle, learn to speak French, were I to enroll in a French class. To have the ability, I, in fact, need to have the relevant psychological mechanism that responds to French utterances.

thus necessarily involves an ability to respond to moral evidence.⁴⁹ The capacity of moral understanding is, on this account, analogous to other epistemic capacities, such as perceptual capacities, for example, the capacity for color vision. One way to characterize this capacity is as the ability to achieve color knowledge via vision when presented with colored objects in various viewing conditions.⁵⁰ Similarly, we can characterize the capacity of perfect pitch as the ability to identify – i.e. know – which pitch is being played when she is aurally presented with a sound. Like these two, moral understanding is a capacity to achieve knowledge when presented with the right evidence. But unlike color vision or perfect pitch, the capacity of moral understanding is not tied to a particular sense modality.

The Moral Knowledge Account tells us what the capacity of moral understanding consists in. But it also entails reductionism about instances of moral understanding. This is because we achieve instances of moral understanding by successfully exercising our capacity of moral understanding. If the capacity of moral understanding is the capacity to acquire moral knowledge, then successfully exercising this capacity results in achieving moral knowledge. As I have argued earlier that there are compelling independent reasons to accept reductionism, this is a strength of the Moral Knowledge Account.

We have an initial sketch of the Moral Knowledge Account on the table. I will now draw some comparisons between it and the moral reasoning claim, with an eye towards both filling out its details and highlighting its advantages.

In contrast to the moral reasoning claim, the Moral Knowledge Account allows for many ways in which an agent can achieve instances of moral understanding because there are many ways in which we can achieve moral knowledge. One such way is moral reasoning. In some situations, there may be no way around moral reasoning to gain moral knowledge. For example, when you face a novel moral scenario, requiring you to weigh difficult factual and conflicting moral considerations. How many refugees a community should accept? Should we allow parents to freely choose schools for their children? Should we give the sole available dose of a hitherto untested Ebola drug to an American who contracted the disease while volunteering or to the local doctor? These are difficult moral questions and insofar as we can achieve moral knowledge about what the right thing to do is, it's only by carefully reflecting about the various options. Thus, having the

⁴⁹ Moral understanding is thus, broadly speaking, a discriminatory ability: an ability to discriminate right from wrong. See Williamson, Timothy, *Identity and Discrimination* (Oxford: Blackwell-Wiley, 2013), chapter 1. What exactly moral evidence consists in is an important question that unfortunately goes beyond the scope of this paper. But I am sympathetic to a view on which moral evidence consists of both moral and non-moral facts. See Kieran Setiya, *Knowing Right from Wrong* (Oxford: Oxford University Press, 2012) for a view along these lines.

⁵⁰ See Will Davis, *Colour Vision and Seeing Colors*, *British Journal for the Philosophy of Science* (forthcoming).

ability to engage in moral reasoning will be part of one's capacity of moral understanding.

The Moral Knowledge Account thus accommodates an important insight of the moral reasoning claim: being bad at moral reflection can seriously limit one's capacity of moral understanding. Being an incompetent moral reasoner generally makes one vulnerable to moral mistakes when it comes to certain moral situations, in which we have to make complex moral decisions. Such a limitation is not inevitable, however. An agent may compensate for her impaired moral reasoning by, for example, being particularly affectively attuned to morally relevant features of the situation. And, conversely, an agent with impaired affective responses may compensate for this by developing her capacity for moral reasoning.⁵¹

The Moral Knowledge Account is pluralist – it allows that the capacity of moral understanding can be multiply realized. Plausibly there are many, distinct faculties and cognitive mechanisms by which we acquire moral knowledge: perception, imagination, intuition, our affective responses, and moral reasoning can all be sources of moral knowledge. On the Moral Knowledge Account then, moral understanding is realized by a set of different faculties and cognitive abilities. Agents can have these cognitive abilities to different degrees. And so, what exactly grounds the capacity of moral understanding may vary from agent to agent.

The Moral Knowledge Account explains why first-personal experience is important to moral understanding. We often say that only by seeing something first-hand “we really got it” or “it finally clicked”. What is it that first-hand experience gets us? Actions are right and wrong in virtue of their features. To understand why an action is wrong generally requires you to know what (some of) its wrong-making features are. First-personal experience gives you a richer conception of what those are. Contrast being told that a patch of color is red with looking at it. In both cases you may come to know that it's red. But in the latter case you learn a lot more: you learn that it's red by seeing its precise shade. Similarly, compare being told that prisons are dehumanizing with visiting a prison yourself. In the latter case, you learn a lot more: you come to know that it's dehumanizing by seeing the myriad ways – big and small – in which prisoners are dehumanized. Of course, you can learn more detail by seeking out more detailed testimony. But, setting aside that what we may lack words to express some of what we see, even the most detailed testimonial account cannot rival the richness of the content of our own perception. We can see this play out in Orwell's case. He sees “the mystery, the unspeakable wrongness, of cutting a life short when it is in full tide” by seeing the very particular things that cutting a life short *was*: someone whose “eyes saw the yellow gravel and the grey walls, and his brain still remembered, foresaw, reasoned – reasoned even about puddles,” would cease to be.

⁵¹ Jeannette Kennett, “Autism, empathy and moral agency,” *Philosophical Quarterly* 52 (2008) (2008), 340-357 discusses this specifically in the context of agents with Asperger's syndrome, arguing that such agents can achieve a high degree of moral reliability by engaging in explicit and conscious moral reasoning.

Second, certain first-personal experiences may be essential to gain epistemic access to certain wrong-making features because these include facts about how the action affects the wronged party: what being the victim of this wrong *is like*. It's plausible that you cannot fully appreciate some important aspects of the relevant experience – the shame of a survivor of sexual assault, or the isolation experienced by a victim of domestic violence – unless you yourself have undergone a similar experience yourself.⁵²

First-personal experiences may not just help us to achieve particular instances of moral understanding – they can also broaden our capacity of moral understanding. Again, on the Moral Knowledge Account, this is not surprising. According to Lewis, by having a particular first-personal experience:

...you gain abilities to remember and to imagine. After you taste Vegemite, and you learn what it's like, you can afterward remember the experience you had. By remembering how it once was, you can afterward imagine such an experience. Indeed, even if you eventually forget the occasion itself, you will very likely retain your ability to imagine such an experience. Further, you gain an ability to recognize the same experience if it comes again.⁵³

Plausibly, what goes for the taste of Vegemite also goes for the (metaphorical) taste of injustice, violence, or poverty. Someone who has experienced poverty knows more about economic injustice because she knows what it's like to be on the receiving end of it. Because of this, she can empathize with others in the same situation, she can imagine what it's like to be in their shoes, she recognizes the telltale signs of inequality. Someone who has experienced being the subject of disparaging sexist comments knows what it's like to be on the receiving end of sexism. For this reason, she can put herself in the shoes of others in similar circumstances, she remembers what it was like for her and she can imagine what it's like for them. These abilities may well make her more sensitive to whether a given remark is inappropriate and more attuned to the impact it has on others. Her experience thus expands her moral understanding because it makes her more sensitive to a certain class of moral evidence. It increases the range of circumstances for which she is in a position to know what the right thing to do is; and so, it broadens her capacity to acquire moral knowledge.⁵⁴ Her capacity for moral understanding is expanded, even if she is not able to

⁵² The emphasis here is on *fully*. Testimony from others and your own imagination can, of course, take you a long way towards such knowledge.

⁵³ David Lewis, "What experience teaches," in *Papers in Metaphysics and Epistemology* (Cambridge: Cambridge University Press, 1999), 260-290, at 286. Lewis argues that knowing what it's like just is a matter of having these abilities. But it's plausible that we acquire such abilities by first-personal experience, even if we do not think that they are all there is to knowing what it's like.

⁵⁴ Of course, the reverse is also true. Certain first-order experiences can leave the agent numb, jaded, or cynical – less likely to notice the needs of others – and so limit one's capacity of moral understanding.

articulate her knowledge – even if she cannot explain to someone else why a particular remark is off.

Many writers on understanding have noted that there's a phenomenal aspect to achieving an instance of understanding.⁵⁵ This phenomenology is varied. This is reflected in how we describe our moments of moral insight: we are hit by the realization that we must help, it dawns on us (sometimes painfully) that our remark was inappropriate, we feel that someone else's action was morally wrong, we see that we must tell the truth. The Moral Knowledge Account explains why: imagination, moral perception, emotional responses, and reflection are all ways of achieving moral understanding and all have their own, distinctive phenomenologies.

7. Second-Hand Understanding and the Limits of Testimony

The Moral Knowledge Account allows that you can gain instances of moral understanding by relying on moral testimony: it allows for second-hand understanding. This chimes well with our moral practice. When we look to understand why an action is right or wrong, we often draw on the expertise of others – those with a wealth of experience and knowledge on that question. The Moral Knowledge Account explains why having the capacity of moral understanding is essential for a good moral advisor. We generally look for moral testimony and advice when we are at a loss about what the right thing to do is and we want to find out. We are looking for moral knowledge.⁵⁶ And so, it's natural that we turn to someone who has the ability to know what the right thing to do is presented with the relevant moral evidence: namely the details of our predicament.

But you may worry that moral testimony still presents a challenge for the Moral Knowledge Account. Consider Sara. Sara's ability to arrive at true moral judgments on her own is extremely limited. But she has access to reliable moral advisors, to whom she willingly defers. Doing so, she generally succeeds in acting rightly. She has the capacity to know right from wrong, you may argue. But since it's solely in virtue of her reliable advisors, we would not credit her with moral understanding.⁵⁷

But the Moral Knowledge Account does not suggest otherwise. First, to acquire knowledge by testimony, you need to be able to recognize reliable advisors as such. If you cannot discriminate among reliable and unreliable moral advisors, it's a matter of luck that you managed to pick a reliable one. Had you come across an unreliable one, you would have formed a false belief, based on their misleading testimony. You are in the same epistemic position as the agent who happens to pick out the one accurate

⁵⁵ See, for example, Linda Zagzebski, *Recovering Understanding*.

⁵⁶ I argue for this account of moral advice in Paulina Sliwa, "In Defense of Moral Testimony," *Philosophical Studies* 158 (2012), 175-195.

⁵⁷ I'm grateful to Stephen Grimm and an anonymous referee for raising this objection.

history book from a pile of misleading ones. Such epistemic luck precludes knowledge.⁵⁸

The ability to recognize a reliable moral advisor requires the capacity to acquire moral knowledge based on evidence other than just moral testimony. Without moral competence of your own – at the very least, the ability to recognize kindness, generosity and good sense, when you see it – you cannot judge other people’s moral competence. Unlike in the case of doctors and lawyers, there is no external seal of moral approval that we can rely on.⁵⁹ Second, recognizing competence is not enough. Someone can be competent in a domain and yet a bad source of advice because they are untrustworthy: they lack integrity, they are sloppy, or insincere. But sincerity, integrity, concern for advising others well are all moral attributes. Sensitivity to them involves sensitivity to moral evidence. As Fricker argues:

Epistemic trust incorporates ethical trust, because epistemic trustworthiness incorporates one kind of moral trustworthiness: namely, sincerity. Accordingly, the virtuous hearer’s sensitivity to epistemic saliences will involve a sensitivity to some moral saliences – seeing a speaker in epistemic colour entails seeing them in some moral colour.⁶⁰

Insofar as Sara’s ability to acquire moral knowledge is seriously impaired in general, this may preclude her from gaining moral knowledge by testimony.

On the Moral Knowledge Account then, the ability to acquire moral knowledge by testimony is partly constitutive of one’s capacity of moral understanding. But this does not make the capacity of moral understanding a free-for-all. Having the ability to acquire moral knowledge by testimony makes substantive epistemic demands on the agent. Simply deferring to someone who, as a matter of fact, turns out to be reliable does not suffice.

Lack of moral competence on a particular question limits your ability to gain knowledge by testimony in another way: as we saw earlier, what you can learn from any particular testimonial exchange depends, in part, on what you already know and which concepts you have. Thus, suppose that Sara has never come across the notion of ableism before – she is neither disabled herself, nor knows any people with disabilities, and she has generally never given the issue any thought. A reliable, trustworthy source tells her that a given remark is “ableist” and nothing more. If she accepts the testimony, what has she come to know? It seems that she may well know that the

⁵⁸ Accepting this does not force us to take sides in the debate between reductionists (for example Elizabeth Fricker, “Against Gullibility,” in *Knowing from Words*, eds. Arindam Chakrabarti & Bilal Krishna Matilal (Dordrecht: Kluwer, 1994) and non-reductionists (for example, Charles Coody, *Testimony: A philosophical study* (Oxford: Oxford University Press, 1992) in the philosophy of testimony. This is because, while anti-reductionists maintain that we have a priori warrant for trusting testimony, they also accept that this warrant is defeasible. And so, it’s plausible that to be in a position to acquire knowledge by testimony, you need to be sensitive to at least some of the defeating conditions.

⁵⁹ See Sliwa, In Defense, for more discussion about how to assess moral expertise without running afoul of general skepticism.

⁶⁰ Miranda Fricker, Epistemic Injustice, 76.

sentence “this remark is ableist” expresses a truth. She may also come to know – perhaps, based on the context or the tone of voice in which the remark was made – it expresses some kind of moral criticism, that it’s some way of saying that the remark is morally unacceptable. But to know this is not the same thing as having knowledge of the proposition expressed by the sentence – knowledge *that this remark is ableist*.

To come to know that the remark is ableist based on testimony, you need to have some prior background knowledge, both moral and non-moral, about the kinds of difficulties that agents with disabilities face, as well as some conceptual competence with respect to the concept of ableism. At the same time, you may not be sensitive enough to the relevant class of moral evidence yourself to recognize many of the instances in which the concept applies. And so, relying on someone with greater moral expertise, you acquire second-hand understanding that making such remark was wrong because it is ableist.

On the Moral Knowledge Account the ability to know based on moral testimony is part of our capacity of moral understanding. You can achieve instances of moral understanding by relying on moral testimony that would otherwise be outside your reach – these are instances of second-hand understanding. Moral testimony builds on and extends our existing moral understanding. It does not create it out of thin air.