



## Oxford Studies in Metaethics, Volume 10

Russ Shafer-Landau

Print publication date: 2015

Print ISBN-13: 9780198738695

Published to Oxford Scholarship Online: August 2015

DOI: 10.1093/acprof:oso/9780198738695.001.0001

# Cognitivism about Moral Judgement

Alison Hills

DOI:10.1093/acprof:oso/9780198738695.003.0001

## Abstract and Keywords

What is it to make a moral judgement? There are two standard views, cognitivist and non-cognitivist, plus hybrid options according to which moral judgements have cognitivist and non-cognitivist components. In this context, cognitivism is typically defined as the theory that moral judgements are beliefs. This chapter aims to clarify what it means for a moral judgement to be a belief. It begins by identifying a tension between three claims: cognitivism, an account of belief, and an account of moral judgement. All three claims are plausible, but the question is whether they can be reconciled. In order to do so, the chapter distinguishes between broad and narrow belief. It gives an example of a mental state that is a broad belief but not a narrow belief, a moral 'ulief'. The chapter sets out new definitions of cognitivism and non-cognitivism and draws out some further implications of the argument.

*Keywords:* moral judgement, cognitivism, non-cognitivism, belief, mental state

When someone tells me that what I am proposing to do is immoral, I do not react by asking for his credentials but for his reasons.

(Coady 1992: 71)

## 1.1 Cognitivism and Non-Cognitivism about Moral Judgement

What is it to make a moral judgement?<sup>1</sup> There are two standard views, cognitivist and non-cognitivist, plus hybrid options according to which moral judgements have cognitivist and non-cognitivist components.<sup>2</sup> In this context, cognitivism is typically defined as follows:

**Cognitivism:** moral judgements are beliefs.

By contrast, there are two standard definitions of non-cognitivism. The first is negative, characterizing non-cognitivism as simply the denial of cognitivism:

**Non-cognitivism 1:** moral judgements are not beliefs.

**(p.2)** The second makes a positive claim:

**Non-cognitivism 2:** moral judgements are 'desire-like' mental states (e.g., desires, pro-attitudes, commitments, emotions, intentions, etc.).

This essay is not a defence of cognitivism (or of non-cognitivism). Rather, the aim is to get clearer about what it means for a moral judgement to be a belief.

I begin by setting out, in section 1.2, a tension between three claims: cognitivism, an account of belief, and an account of moral judgement. I think all three claims are plausible, but I will not be defending them here. Rather, my interest is in whether they can be reconciled. In section 1.3, I distinguish between broad and narrow belief. In sections 1.4 and 1.5 I give an example of a mental state that is a broad belief but not a narrow belief, a moral 'ulief'. Section 1.6 addresses some objections to the idea of a moral ulief. In the final section I set out new definitions of cognitivism and non-cognitivism and draw out some further implications of the argument.

## 1.2 Belief, Moral Judgement, and Trusting Testimony

Cognitivism is an appealing theory of moral judgement. Like beliefs, moral judgements appear to describe the world. The semantic expression of moral judgements, like beliefs, is in declarative sentences that are truth-evaluable.

Cognitivism faces difficulties, naturally, when moral judgements do not appear to behave like beliefs. One such is very well known. It is widely believed that moral judgements are 'practical': when you make a moral judgement, *ceteris paribus*, you are motivated to some degree to act on that judgement. But beliefs have no such connection to motivation.

This has a good claim to be the most discussed argument in metaethics. I am not convinced by it, but I will not be adding to the already vast literature. Instead, I am going to discuss a different problem for cognitivism, a tension between it and two further claims, one about the nature of beliefs, and one about moral judgement.

The claim about beliefs is: beliefs should be responsive to evidence. This might be a descriptive claim—that beliefs typically and characteristically do respond to evidence—but I will mostly discuss the normative claim—that they should.

Of course, beliefs are not always formed as they should be; we are careless, or self-deceived, or victims of wishful thinking. But if it's brought to your attention that you have formed a belief that is not properly supported **(p.3)** by evidence, you should suspend judgement or form a new belief that fits better with your evidence.

One of our most important sources of evidence about the world is testimony: under the right circumstances, testimony that p is good evidence that p.

For example, suppose you are having a discussion about tax, and someone says 'an increase in sales tax of 1% will raise X millions for the government'. You might check her credentials: is she really knowledgeable about tax? Has she thought about the matter carefully? And if her credentials are fine, you should trust what she says, because by doing so you can gain knowledge.

This is an example of what I call 'pure' testimony, that is, testimony that p without any indication of the reasons why p might be true.

Of course, things are not always this straightforward. First, you should not trust testimony when the speaker is unreliable, trying to deceive, or has simply made a mistake. Here, testimony is not good evidence, so beliefs should not be responsive to it.

Secondly, testimony is only one piece of evidence.

Suppose that a speaker says that p. But you have lots of evidence that not-p. Her evidence is *outweighed*.

Suppose she says that p, but you have lots of other, very strong evidence for p. Her evidence is *swamped*. Though you do form the belief that p, the belief is responsive to other, stronger evidence and is responsive only a little, if at all, to her testimony.

But even in these cases, your beliefs are responsive to pure testimony insofar as you treat it as evidence in its own right that is assessed along with the other evidence that you have regarding p, and which plays a role in determining whether or not you believe that p.

In short: beliefs should be (and are) characteristically responsive to testimony that p, provided that it is regarded as good evidence that p.

The problem for cognitivism arises when this claim about belief is combined with a claim about moral judgements, namely that moral judgements should not be responsive to pure moral testimony.

Suppose that in your tax conversation, someone says: ‘an increase in sales tax is morally wrong’. Should you check her moral credentials, make sure that she was knowledgeable about tax, and then trust her moral judgement? Some people say: yes you should. Doing so will get you moral knowledge.<sup>3</sup> **(p.4)** If this is right, moral judgements are just like beliefs, in this regard at least, and there is no problem for cognitivism.

But other people say no.<sup>4</sup> Consider the quote from Coady with which I began which says that you should not trust pure moral testimony, instead you should ask for the reasons. Why is an increase in sales tax wrong? You are looking for an explanation that you yourself can follow and accept or reject. That is, you are not ignoring this testimony, but you are not putting your trust in it either. You are assessing the case for the moral conclusion yourself.<sup>5</sup>

When we first make moral judgements as children, we are strongly influenced by those around us. But many people regard it as an ideal that, at a certain point, people make up their own minds about morality without relying on the authority of others. What kind of ideal is this? Moral? Presumably. Epistemic? That isn’t so clear, and I will return to this question later. It is striking that moral philosophers who are in deep disagreement about ethics—including Kant and Bernard Williams—all say that trusting moral testimony is not the best way of making moral judgements.<sup>6</sup> Aristotle **(p.5)** too suggests that a virtuous person will have good judgement (*phronesis*) and should use it to decide what to do.<sup>7</sup> I assume (and will later defend the assumption) that it is at least not epistemically irrational to form moral judgements in this way.

It is important to be clear about what this means, because of course, testimony can still be important to these people when they are thinking about morality. Their judgements should be responsive to testimony about relevant non-moral matters.<sup>8</sup> And these judgements should be responsive to moral reasons and moral argument. But *ceteris paribus*, they should not be responsive to pure moral testimony, that is, to a speaker simply asserting that X is morally wrong.<sup>9</sup>

Why shouldn’t moral judgements be responsive to pure moral testimony?

One possibility is that there is no good evidence for any moral judgements. Perhaps there are no moral truths; perhaps there are, but we have no good access to them. Assessing this claim properly is beyond the scope of this chapter. But in any case, this cannot explain why moral judgements should be insensitive to pure moral testimony, but sensitive to other kinds of evidence (including non-moral facts, moral arguments, thought experiments, and so on).

A second possibility is that though there is some good evidence for our moral judgements, moral testimony is not among it. Other people’s moral judgements are unreliable, so we should not place our trust in them.

This argument might be framed in terms of responsibility for our beliefs. When forming beliefs about very important matters—as moral questions frequently are—the responsible thing to do is to make up your own mind, rather than to put your trust in others.

Of course, it would not be responsible to trust unreliable testimony, and much, perhaps most, moral testimony is unreliable. It is difficult to find someone whose views you can trust in ethics. There are wide divergences in opinion and sometimes people's own interests can be at stake, so that their judgements may be biased. This is certainly an important factor in explaining why we should not trust moral testimony.

But there is no reason to think that moral testimony is always unreliable. It is implausible that other people's moral judgements are, quite generally, much less reliable than one's own. And then why not trust moral testimony that is, and that you take to be, reliable? Surely this would actually be the **(p.6)** most responsible way of forming moral beliefs: more responsible than using your own judgement, when that too can go wrong.<sup>10</sup>

A third possibility is that our moral judgements actually should be sensitive to evidence from pure moral testimony, but that evidence is either outweighed or swamped by other evidence that we have. It may seem so obviously true that it is wrong to torture innocent children that testimony either for or against the practice can safely be ignored.

But again, it is not plausible that other evidence I have always outweighs or swamps moral testimony. Not all moral questions have obvious answers, and other people may well be better at making moral judgements than I am.

More importantly, when I make a moral judgement, I do not always think: 'this person is unreliable; I can and should ignore what she says because I am more likely to get the answer right that way'. Instead I think: 'this is the sort of issue about which I should make up my own mind. I will listen to her arguments but in the end I need to come to my own conclusion rather than put my trust in her'.

In other words, if I can make up my own mind properly about a moral question, I should: in these circumstances, trusting pure moral testimony is 'the wrong kind of reason' for forming a moral judgement, quite aside from any concerns about reliability.

Now we can see the tension for cognitivism. Anyone who accepts this claim about moral judgement thinks that pure testimony can be the right kind of reason for a belief, and at the same time 'the wrong kind of reason' for forming a moral judgement.

An easy way out of the problem for a cognitivist would be to deny the claim about belief (it should be responsive to evidence) or the claim about moral judgement and moral testimony. But what if, like me, you are persuaded by both of them? Can you still be a cognitivist about moral judgement? How can the tension between those two claims and cognitivism be resolved?

One possible move for a cognitivist is to accept a simple form of moral subjectivism. If moral truths were dependent on an individual's attitudes, there would be little point in trusting someone else's moral judgements because what was true for them on the basis of their attitudes might not be true for you. Your own attitudes might be different from theirs, and you (p.7) might be the best placed to work out what they are. So it might be best to make moral judgements using only your own judgement.

Simple subjectivism would be a way of reconciling cognitivism with the two further claims. Unfortunately this very simple form of subjectivism leads to an extreme kind of moral relativism: moral truths are relative to each person. This is not very appealing. On the other hand, a more sophisticated kind of subjectivism that does not lead straight to relativism cannot explain why you should not trust moral testimony. So it does not seem that cognitivists can solve their problem satisfactorily by appealing to subjectivism.

There have been some other attempts to explain why we might be reluctant to trust moral testimony.

### (1) The importance of autonomy.

Autonomous judgements are self-given or self-made. If it were important that moral judgements were self-given, this could explain why you should use your own judgement rather than trusting moral testimony, because putting your trust in someone else's judgement is not to make a judgement yourself.

### (2) Motivation.

Moral judgements tend to motivate people to act. They are essentially practical. It may be that taking a moral judgement from someone else, on trust, tends not to motivate. Hence, there is reason to make your own moral judgements instead.

I am not convinced that either of these explanations will ultimately be successful (McGrath and Hopkins argue convincingly that they will not).<sup>11</sup> The second is problematic for cognitivism because it depends on an aspect of moral judgements that is desire-like rather than belief-like. The first requires further explanation because it is not clear why some beliefs (but not others) should be formed 'autonomously', that is, without trusting testimony.

In the remainder of the chapter, I turn to a different kind of solution. I begin with a distinction between two types of belief.

---

**(p.8)** 1.3 Broad and Narrow Belief

It has been acknowledged for some time that the term 'desire' can be used in a broad or a narrow sense. In the broad sense, 'desire' means a non-cognitive 'desire-like' attitude, sometimes called a pro-attitude. A desire in the narrow sense is one of those attitudes, but there are others. For instance, here is Davidson characterizing 'pro-attitude' as including:

Desires, wantings, urges, promptings, and a great variety of moral views, aesthetic principles, economic prejudices, social conventions, and public and private goals and values.

(Davidson 1980: 4)<sup>12</sup>

And here is Bernard Williams characterizing a 'subjective motivational set' or S:

I have discussed S primarily in terms of desires, and this term can be used, formally, for all elements in S. But this terminology may make one forget that S can contain such things as dispositions of evaluation, patterns of emotional reaction, personal loyalties, and various projects, as they may be abstractly called, embodying commitments of the agent.

(Williams 1981: 105)

Desires in the narrow sense are a subset of desires in the broad sense.<sup>13</sup> A very similar distinction can be made between a broad and narrow sense of belief (what I will call broad belief and narrow belief or n-belief).<sup>14</sup>

**(p.9)** A broad belief is a representational state, a propositional attitude. It has a 'mind to world' direction of fit. It has a characteristic functional role, being typically formed in response to evidence, and taking part in reasoning and in action. It may also have a characteristic phenomenology:

This what-it-is-likeness typically involves: (1) psychologically "coming down" on some issue, in a way that (2) classifies (sometimes spontaneously) some "object" of focus as falling under some category, where one's classificatory coming down is experienced (3) as involuntary, (4) as a cognitive response to some sort of consideration that is experienced (perhaps peripherally in consciousness) as being a sufficient reason for categorizing as one does, and (5) as a judgement that is apt for assertion and hence is naturally expressible in public language by a sentence in the declarative mood.

(Horgan and Timmons 2006: 263)

Broad beliefs can be divided into sub-kinds. One sub-kind is a narrow belief. Narrow beliefs are propositional attitudes that aim at truth in a distinctive way: they aim at knowledge.<sup>15</sup> These metaphors of 'aiming' are notoriously obscure.

What I mean here is: a narrow belief that *p* should be responsive to evidence regarding *p* in such a way that in favourable circumstances the result will be knowledge that *p*.<sup>16</sup> I am not of course suggesting that we consciously try to respond to evidence in this way, nor that we deliberately aim at knowledge. Most of us do not. And I am not suggesting that beliefs are always perfectly responsive to evidence. Sometimes we fail to recognize evidence, sometimes we do not form beliefs properly in response to it. Moreover, there may be other influences on belief (including desire). But nevertheless narrow beliefs should be formed in response to evidence, I suggest, in such a way that in favourable circumstances the result will be knowledge and when they are not, we typically regard this as a mistake which we try to correct.

In favourable circumstances, trusting pure testimony from a reliable source will give you knowledge. It follows that narrow beliefs that aim at knowledge should be (and characteristically are) responsive to this kind of evidence.

**(p.10)** Moral judgements that should not respond to pure moral testimony do not have all the characteristics of narrow beliefs. So anyone who accepts that should reject a version of cognitivism of the following form:

**Cognitivism (narrow):** moral judgements are narrow beliefs.

But recall that *n*-beliefs are a sub-kind of broad beliefs. There may be other sub-kinds: types of broad belief that are not narrow beliefs.

In the remainder of this chapter, I introduce a mental state which I call a moral 'ulief'. Moral uliefs are propositional attitudes which represent the world. They are most naturally expressed in declarative sentences which are truth-evaluable. They have the 'direction of fit' that they 'aim to fit the world'. They are therefore cognitive states (broad beliefs). But they are not narrow beliefs: they do not aim at knowledge. The best way to think of the difference between narrow beliefs and uliefs is that as moral beliefs stand to moral knowledge, so moral uliefs stand to what I call *exercised moral understanding*. But this clearly needs explanation, and I will begin with the idea of moral understanding.

#### 1.4 Moral Understanding

By moral understanding, I mean understanding why *p* where *p* is some moral proposition (such as, killing is morally wrong, giving money to charity is the thing to do). I have introduced and discussed my conception of moral understanding elsewhere, so I shall simply sketch briefly its main features in this section.<sup>17</sup>

Suppose that you understand why giving money to charity is good. Then you judge that it is good, and you grasp why. What is it to grasp why giving money to charity is good? It involves having a set of abilities. For instance, you can give an explanation of why it is good: because doing so helps people who are in great

need. You can give explanations and draw conclusions about similar cases (e.g., when there are other ways of helping people, when the charity in question is not very efficient, and so on).

In some ways, moral understanding is quite similar to moral knowledge. As I conceive of it, it is factive (you cannot understand why p, if p is false or if you are wrong in judging that q is why p). And it is not transparent (you may think you understand why p when you do not, and vice versa). But the distinctive feature of moral understanding is this grasp or set of abilities, which goes beyond what is required to *know that* some action is right, and **(p.11)** even what is needed to *know why* it is right.<sup>18</sup> If you understand why p (and q is why p), then you judge that p and that q is why p and in the right sort of circumstances you can successfully:

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

To understand why p, you have to have the abilities (i)–(vi) to at least some extent.<sup>19</sup>

It is clearly possible to have knowledge that p and even knowledge why p without understanding why p, because you lack the 'grasp' distinctive of understanding. I think that it is also possible to understand why p without knowing why p (or knowing that p). For there may be defeaters present for knowledge, that are not defeaters for understanding.

For instance, your judgement that p could depend on your judgement that q, and you might be lucky in correctly judging that q. For instance, suppose that you read in an otherwise very inaccurate textbook that Pol Pot was responsible for the deaths of millions. You correctly judge that he was evil. You could not know that he was evil (since your judgement was made on an unreliable basis) but you could understand why he was.<sup>20</sup>

**(p.12)** Understanding why p and knowledge why p are thus two separate states. One way of conceiving of the difference uses a Platonic metaphor. Knowledge requires a true judgement, and the judgement has to be 'tied down'. It is controversial exactly what this tie is. Perhaps the true judgement must 'track the truth' or be 'safe' or not be 'lucky'. Understanding why p also requires a true judgement, and also requires that the judgement be tied down. This 'tie' consists in your having a set of abilities to make judgements and give

explanations in this and in similar cases. Exercising understanding requires an even firmer tie: that you have made your judgement by exercising your abilities. Notice that the kind of anchor required for understanding is quite different from that needed for knowledge. You can have either kind of 'tie' without the other: knowledge without understanding or understanding without knowledge.

How can you acquire moral understanding? You need to develop the key abilities, to explain why actions are right and to make judgements about what is right based on the right reasons. Trusting when you are told that this or that action is right will usually not give you the grasp of the reasons why that is essential to understanding.

Why acquire moral understanding in the first place? Because it is important to use the abilities characteristic of moral understanding to decide what to do.

To see why, we need to think about moral virtue. Moral virtue is widely agreed to be a kind of responsiveness to moral reasons, usually explained as responsiveness through action and through feelings, emotions, and desires. But it is also possible to respond appropriately or inappropriately to moral reasons through your cognitive attitudes.

What is appropriate responsiveness to moral reasons? I suggest it is: forming a moral judgement (e.g., this action is morally good) through your sensitivity to the reasons why it is true (e.g., because this action helps people who are in great need). In other words, it is a matter of making moral judgements by exercising your moral understanding.<sup>21</sup>

To have virtue, you need to decide what to do by exercising your moral understanding (i.e., by using some of the abilities essential to moral understanding).<sup>22</sup> In Aristotelian terms, you need *phronesis*. This has serious implications for moral testimony.

**(p.13)** Suppose that you have a friend who is very seriously ill and she asks you to help her kill herself. You are not sure what to do. You ask your parish priest, for whom you rightly have the highest regard. If you put your trust in what he says, you will know what to do.

Suppose that you do. Now both you and he have moral knowledge. But only he has responded directly to moral reasons: if he has given you pure moral testimony, you are not aware of those reasons at all. You have just done what he told you to do. This is not an ideal of moral virtue. You may be well motivated and well meaning but you fall short of full moral virtue, even if you do the right thing knowing that it is right.

Similarly, for your action to be fully morally worthy, you must do the right thing for the right reasons, that is, for the reasons that make that action right.

Suppose that you trust your priest when he says that it is acceptable for you to help your friend and act on that basis. Your action is based on his pure moral testimony. It is not based on the reasons that make it right—it could not be—because you have no idea what they are and you are not responsive to them in any way. Again, you may be well motivated and do the right action, but that is not enough for your action to be fully morally worthy.<sup>23</sup>

Exercising your moral understanding to make moral judgements is an ideal. Other things being equal, you should try to live up to this ideal. But other things are not always equal. First, trusting testimony has a crucial role in moral education. To acquire moral understanding, you need to start with some (preferably true) moral claims, and you may have to take these on trust.

Secondly, there may be some reason why you cannot gain or cannot use your moral understanding. Perhaps you lack certain kinds of experience and these are essential to good judgement. Perhaps you are just not very good at weighing up values. Hopefully you will be able to get good advice that will help you to reach a good judgement by explaining why some action is right and others wrong. But perhaps you are not in a position to get or to appreciate advice. It may be better for you to trust moral testimony instead. Again: this **(p.14)** is a sign that you are not fully morally virtuous.<sup>24</sup> Ideally, you would make up your own mind on the basis of reasons that you yourself appreciate.

### 1.5 Moral Uliefs

Suppose that an increase in sales tax has been proposed and you are wondering whether this policy is just. You recognize that raising this tax will force the worse off to pay a higher rate of tax than those who are better off. You judge that it is unjust on this basis. To put it schematically, you are wondering whether p, whilst you recognize that q. Exercising your moral understanding, you judge that p and that q is why p.

Let us focus on your judgement that p, that a rise in sales tax is unjust. What kind of mental state have you now formed?

The state is plainly a propositional attitude. It can play a role in action. For instance, you might oppose the policy of an increase in sales tax, because you judge it to be unjust. And you might use the judgement in further reasoning (for example, since further tax income is needed, but a sales tax increase is unjust, a different way of raising tax must be found). The state is responsive to certain kinds of evidence regarding p, for example, arguments, counterexamples, thought-experiments and so on.

At this point, we can describe your judgement as a cognitive state, a belief in the broad sense. But is it a belief (narrow sense)? It is characteristic of narrow beliefs to aim at knowledge. But we have seen that knowledge is not the only epistemic aim. Instead, you might make moral judgements as the virtuous

person does, by exercising her moral understanding. I call a state aimed at this goal, a moral 'ulief'.

You need not be consciously aware of trying to exercise moral understanding to be forming uliefs. Nevertheless uliefs should be formed in response to evidence, in such a way that in favourable circumstances the result will be exercised moral understanding.

If you exercise your moral understanding, you respond to certain kinds of evidence only—evidence regarding the reasons why a rise in sales tax is unjust. But testimony of the form: 'a rise in sales tax is unjust' is not a reason **(p.15)** why that proposition is true. So if you are forming uliefs, you will typically not put your trust in testimony of that form, however reliable or expert you regard the speaker.

In other words, if your moral judgement that a rise in sales tax is unjust is sensitive to evidence including moral testimony, it is a (narrow) belief. If it is sensitive to other kinds of evidence, but not moral testimony (or similar), it is a moral ulief.

There are a host of well-known problems for any view that moral judgements are not beliefs. These include requirements to explain how such judgements are truth-apt and subject to logical relations and how they can be a proper focus for disagreement. These are all extremely demanding challenges for non-cognitivists to meet. Are they equally difficult for a 'moral ulief' account of moral judgement?

Not at all. Moral uliefs are a kind of broad belief, therefore they are factive states that 'aim at the truth'. Moral uliefs, like narrow beliefs, will typically be expressed by declarative sentences with descriptive content that can be assessed as true and false. For instance:

Sales tax increases are unjust.

Sales tax increases are unjust because they disproportionately affect the worst off.

Since truth is as important for uliefs as beliefs, uliefs can and do play the same role in reasoning as beliefs. You can derive further uliefs or beliefs from combinations of your current uliefs and beliefs in familiar ways. For example, you might think though the following argument:

1. Rises in tax which disproportionately affect the worse off are unjust.
2. A rise in sales tax is a rise in tax which disproportionately affects the worse off.

Therefore

3. A rise in sales tax is unjust.

A piece of mental reasoning like this may be constituted by narrow beliefs—that is, you may *n-believe* the propositions expressed by (1) and (2) and draw the appropriate conclusion, forming a narrow belief that a rise in sales tax is unjust. But here is another possibility. You may *n-believe* that this is a tax rise that disproportionately affects the worse off, but have a moral *ulief* that rises in tax which disproportionately affect the worse off are unjust (e.g., because you grasp that a just system of taxes is sensitive to need and the ability to pay). And you may as a consequence draw the appropriate conclusion, forming a moral *ulief* (that this rise in tax is unjust) rather than an *n-belief*. Moral uliefs can play a role in reasoning and in valid arguments **(p.16)** because the ‘logic’ of moral uliefs is sufficiently similar to the ‘logic’ of narrow beliefs. And finally, if you form the moral ulief that killing the innocent is always wrong, whilst I form the moral ulief that it is not, it is perfectly clear that we disagree with each other. Our judgements cannot both be true, and anyone who ‘ulieved’ both propositions would have inconsistent uliefs. Just as there is something *prima facie* wrong with having inconsistent *n-beliefs*, there is something *prima facie* wrong with having inconsistent uliefs (or indeed in having an inconsistent set of uliefs and *n-beliefs*).

The role of belief, particularly moral beliefs, in action is of course a matter of dispute. On some views, beliefs cause action only in combination with a desire. According to other theories, beliefs can sometimes cause action on their own, without the need of an additional desire. I will not take a stand on this controversy. I assume that whatever role moral *n-beliefs* play in causing action can also be played by moral uliefs.

I have introduced the idea of moral uliefs by describing moral understanding, its use in virtue and morally worthy action. In the next section, I want to consider some objections that might be raised to this account of moral uliefs as belief-like mental states.

## 1.6 Objections

### 1.6.1 Is a Uliel a Kind of Cognitive State At All?

I will take it that a cognitive state is a broad belief. In section 3, broad beliefs were defined in terms of a set of properties. It is a propositional attitude. It has a ‘mind to world’ direction of fit. It has a set of characteristic functions—being typically formed in response to evidence, taking part in reasoning and in action. It may have a characteristic phenomenology:

This what-it-is-likeness typically involves: (1) psychologically “coming down” on some issue, in a way that (2) classifies (sometimes spontaneously) some “object” of focus as falling under some category,

where one's classificatory coming down is experienced (3) as involuntary, (4) as a cognitive response to some sort of consideration that is experienced (perhaps peripherally in consciousness) as being a sufficient reason for categorizing as one does, and (5) as a judgement that is apt for assertion and hence is naturally expressible in public language by a sentence in the declarative mood.

(Horgan and Timmons 2006: 263)

I think it is clear that moral uliefs will have all these features. Uliefs aim to fit the world, though in a slightly different way from n-beliefs: I have characterized as n-beliefs aiming at knowledge whilst moral uliefs aim at exercised moral understanding. Moral uliefs have the functional role of a **(p.17)** cognitive state: they can play the same role in argument as n-beliefs and can be used in the same mental inferences (aside from any specifically relating to testimony or deference to experts or similar). They do not have exactly the same functional role with respect to action as narrow beliefs, because they play a crucial role in fully morally worthy action and in full moral virtue. But nevertheless, their causal contribution to action is in other respects similar.

Finally, moral uliefs have the phenomenology of a cognitive state. A moral ulief is formed on the basis of your grasp (or your attempted grasp at least) of the reasons why it is true. You cannot just choose to judge whether a rise in sales tax is unjust or not, once you recognize (what seem to you to be) the reasons why it is unjust. This is no doubt experienced as a 'coming down' in classifying the rise in sales tax ('the object') as falling into a particular category (unjust policies). It is naturally expressed in an assertion ('A rise in sales tax is unjust') by a declarative sentence.

I leave it open whether there are cognitive states other than narrow beliefs and uliefs, though I think it likely that there are.<sup>25</sup> But in any case, there is no doubt about it: uliefs are a type of cognitive state.

#### 1.6.2 Why Distinguish Uliefs and N-Beliefs? Why Not Call Them All Beliefs?

No one has until now suggested that there is such a thing as a ulief, a cognitive state that is not a (narrow) belief. Surely, if there were any such states, we would know about them? But though many people who make moral judgements describe them as beliefs ('I believe that a rise in sales tax is morally wrong'), absolutely no one has ever before talked about her uliefs.

In addition, isn't proliferating the number of different cognitive states multiplying entities beyond necessity? Why not just say that there is one fundamental type of cognitive state: a belief? It is extravagant to posit two different states to do basically the same thing.

With regard to the first objection: it is legitimate to use the term 'belief' as a general term for all cognitive mental states (I have done so here, though for clarity I have called them broad beliefs). It is certainly true that no one talks about uliefs. Does this matter? Only if we are some sort of authority on our own mental processes, that these are transparent to us. But obviously they are not. A failure to distinguish between two cognitive states is hardly unexpected.

**(p.18)** With regard to the second objection, it is true that the criteria for individuating types of mental state are not clear. Nevertheless, I have argued that there are some good reasons for distinguishing between different 'belief-like' cognitive states, whether these are considered as different *types* of mental state or different *sub-kinds* of broad belief. Some of these reasons are theoretical. I have argued that knowledge and (exercised) moral understanding are different states, and that you can have one without the other. Narrow belief 'aims' at knowledge, hence it is responsive to all evidence. Moral ulief 'aims' at exercised moral understanding, hence is responsive only to certain kinds of evidence.

Secondly, there are practical reasons for distinguishing n-beliefs and uliefs. Moral virtue is typically a matter of forming and acting on moral uliefs, rather than n-beliefs.<sup>26</sup>

These are good reasons for concluding that n-beliefs and uliefs are two different mental states. Nevertheless, there are costs to doing so. The major difficulty consists in explaining the relationship between the two: when you have a ulief that p do you (Sometimes? Always? Never?) have a narrow belief that p as well? Can you have a ulief that p and a belief that not-p at the same time? What are the inferential relations between the two? A full account of this relationship is beyond the scope of this chapter, but in the following section, I will make a start, hoping to say enough to vindicate drawing the distinction between uliefs and n-beliefs.

### 1.6.3 What Is the Relationship between Moral Uliefs and Moral N-Beliefs?

According to some people, all moral judgements are narrow beliefs: all should be responsive to pure moral testimony. Even those who deny this, and think that not all moral judgements should be so responsive, nevertheless typically think that at least some of them should: some moral judgements are narrow beliefs.

Suppose that you begin with a narrow moral belief, formed on the basis of testimony from your family or friends. But you begin to reflect, using your moral understanding, on why it is true—and why it might not be. Suppose that you change your mind. You n-believed that abortion was always wrong and now you form the moral judgement that it is not. Your n-belief has been replaced by a ulief (if it had not, you would have believed (in a broad sense) a contradiction).

But what if you reaffirm your prior view? Is your n-belief supplemented or replaced with a ulief?

**(p.19)** I suspect that typically it is replaced, that is, you suspend your n-belief and replace it with a ulief that p. To form a ulief at all, you need to base it on the exercise of your moral understanding, and I think that that is difficult to do (though perhaps not impossible) when you still have an n-belief that p. For you already have the answer to your question: *is p true?* Instead you need to regard the question as reopened, by suspending judgement on whether p.

Of course at this stage you might ulieve that p on the basis of your moral understanding and n-believe that p too. But there are a couple of reasons why the n-belief might remain suspended. Once you have formed a ulief, you have reached a settled opinion that p, you have no need of an n-belief that p as well. And secondly, it makes sense for us to replace n-beliefs with uliefs rather than to supplement them, for if the argument sketched here is right, it is morally important that it is the ulief and not the narrow belief that plays a role in further action, and if there is no n-belief, there is no chance that it will play the role instead.

This is not to say that it is impossible to have both an n-belief and a ulief that p at the same time. Indeed, I also think it possible to n-believe that p and ulieve that not-p at the same time. Perhaps your moral views are not completely consistent, maybe you are in the process of a moral 'conversion'. Consider the following example.

#### **1.6.3.1 The Milgram Experiments**

In the famous Milgram experiments, participants were asked to 'shock' a victim to 450 volts (starting with a shock of 15 volts and increasing at 15 volt intervals) despite his screams of pain and demands to be released. The shocks (which were not real) were apparently applied in the course of an experiment purporting to be about learning. If the subjects raised questions about the experiment, the experimenter, who was present in the room, ordered them to continue (saying, for instance 'please continue', 'the experiment requires that you continue', 'you must go on'). Before the experiment, most people say that they would refuse to participate, but in fact 65 per cent of the subjects continued to shock the victim until they reached 450 volts.

What is going on in the minds of participants in the experiment? This is a highly controversial matter, about which there are different opinions. Here is one of them: the subjects treated the experimenter as a moral expert:

He is the authority—not (or not just) in the sense of the person in charge but in the sense of the person who knows what is to be done in the experimental situation. The subjects, after all, are strangers to the world of the laboratory or, better, visitors to this world. They are looking for clues

about how to act, and the experimenter provides them ... The experimenter shows the subjects a vision of the moral order. **(p.20)** He “tells” them that the right thing to do here is to shock the victim. He does this not by arguing explicitly that it is morally correct to continue shocking but simply by ordering them in the most matter-of-fact way to continue. In so doing, he shows them that he (a seemingly reasonable, smart, competent fellow) takes it to be morally appropriate to do so.

(Sabini and Silver 2005: 550–1)<sup>27</sup>

According to Sabini and Silver, the subjects *n-believe* that it is appropriate (perhaps even required) to continue to shock their victim. This narrow belief is based on their taking the experimenter to think so too, on the basis of his behaviour, and on their taking him to be an expert on the appropriate way to behave in his laboratory. Of course they have no idea why the experiment might be acceptable. And moreover they do not wholeheartedly judge that it is right, even during the course of the experiment. They hesitate, protest, feel under stress, and in a few cases refuse to participate further.

One possibility is that there is a conflict between their cognitive states (their belief that the experiment is acceptable) and their non-cognitive attitudes (their desire not to harm). Another possibility, however, is that there is a conflict in their cognitive states. They *n-believe* that the experiment is morally acceptable, but they also *ulieve* that it is not. This *ulief* is based on their hearing the screams and cries from the subject in the next room: they appear to be causing considerable suffering and possibly injury, for no obvious gain. They naturally form the *ulief* that continuing to participate is morally wrong. This cognitive conflict causes their confused response: they carry on with the experiment thanks to their *n-belief* but their *ulief* causes them to hesitate; in a few people, the *ulief* wins out and they refuse to go on.

Prior to the experiment, almost everyone claims that they would not participate. Why do they say this? Obviously, they think that in the course of the experiment, they would form a *ulief* that participating is wrong, and that they would and should act on that *ulief*. But, as Sabini and Silver put it, in the setting of the experiment, most participants lose their ‘moral rudders’ or their ‘moral compass’, that is, they fail to act on the basis of their own moral understanding. There are a number of interfering factors in the experiment, some motivational (a desire to obey authority, a wish to avoid the embarrassment of making a scene) and some cognitive (a belief that the experiment must be morally acceptable, because the experimenter appears to think that it is). The two also combine: a strong desire not to disrupt the experiment and disobey the authority figure makes the *n-belief* that participating is morally acceptable much more attractive than it would otherwise **(p.21)** be. It is hard—much harder than we

predict—to form a clear ulief that participating in the experiment is wrong, and it is harder than we predict to act on that ulief.<sup>28</sup>

According to this interpretation, the Milgram participants have a clash of cognitive attitudes. Even if you agree that that is possible, you might insist that these clashes are straightforward conflicts of beliefs. But I think that there are clear advantages in regarding them as involving both uliefs and n-beliefs, which are also reasons for making the distinction among cognitive attitudes in the first place.

Interpreting these situations as ulief/narrow belief conflicts emphasizes the different grounds of the cognitive states in each case—whether it is responsive to the authority of other people’s judgements or not. This distinction is worth making, linked as it is to the differences between knowledge and moral understanding. But I have also argued that it is morally important. A virtuous person would be disposed to use their own judgement and would respond to moral reasons appropriately in this sort of situation.<sup>29</sup>

### 1.7 Cognitivism Redefined

Cognitivists about moral judgement have typically claimed that moral judgements are beliefs and the view has often been defined in those terms. I have argued that this is a confusing definition. There are cognitive states (broad beliefs) that are not narrow beliefs, so it would be better to define cognitivism about moral judgement as the following view:

**Cognitivism:** moral judgments are cognitive states (beliefs in the broad sense, i.e., they may be narrow beliefs or moral uliefs or ...).

**(p.22)** And correspondingly, non-cognitivism should be defined either as the negative view:

**Non-cognitivism 1:** moral judgements are not cognitive states (broad beliefs).

Or, if you prefer a ‘positive’ characterization:

**Non-cognitivism 2:** moral judgements are ‘desire-like’ mental states (e.g., pro-attitudes, commitments, etc.).

Since both n-beliefs and uliefs are cognitive states, a cognitivist who thinks ideal moral judgements are not made on the basis of pure moral testimony could claim that whilst some moral judgements are narrow beliefs, others are moral uliefs (a cognitivist who disagreed could maintain that moral judgements are all narrow beliefs). Hybrid theorists, who think that moral judgements are a mixture of cognitive and non-cognitive states could state their theory in terms of uliefs as well as narrow beliefs.

The main purpose of this chapter is to clarify one option open to cognitivists in the face of a problem about pure moral testimony. Of course, this does not establish that cognitivism is correct. In particular, nothing I have said here resolves the ‘practicality’ problem for cognitivism, for moral judgements appear to have a connection to motivation which cognitive states do not. I am not convinced by this argument, but I will not make the case against it here.<sup>30</sup>

Instead, I want to mention very briefly two wider implications of the argument. First, though I have discussed only moral understanding here, I think that understanding differs from knowledge in other fields too. For instance, aesthetic understanding, unlike aesthetic knowledge, requires a grasp of why a work of art is aesthetically valuable; philosophical understanding, unlike philosophical knowledge, requires a grasp of why a philosophical thesis is true. We typically form judgements on these subject matters by making up our own mind, rather than trusting testimony, that is, we form uliefs, not narrow beliefs. Perhaps it is possible to form a ulief on any subject matter on which you can form a narrow belief: I suspect that it is. I hope to say more about this in future work.

**(p.23)** Secondly, I raised a question earlier about whether the ‘ideal’ of forming moral judgements by making up your own mind, rather than trusting moral testimony was a moral or an epistemic idea. I have argued so far that it is a moral ideal, because a fully virtuous agent would form moral judgements that way. But I want to now suggest that it is an epistemic ideal too.

The study of the relationship between the world and our representations of it has traditionally been called ‘epistemology’, and recently this has been understood as being—by definition—the study of knowledge. But I suggest that that is too narrow. Knowledge and justified (narrow) belief is only a subset of that broader subject matter, which includes understanding and uliefs as well. And we have very good reason to think that the epistemic rationality and epistemic justification for understanding, particularly with regard to testimony and related matters (notably, disagreements), will be rather different.

Suppose that you ignore a piece of moral testimony from someone whom you correctly regard as reliable. Could this be epistemically rational? If epistemic rationality is a matter of gaining relevant knowledge, presumably not, because (*ceteris paribus*) you could gain knowledge by doing so. It is true that you might have good moral reasons to ignore testimony—it is what a virtuous person would do, after all—but these are what are sometimes called the ‘wrong kind of reasons’. They are practical reasons, not epistemic ones.

I suggest that this argument is not quite right. ‘Epistemic’ rationality and justification for moral uliefs should relate to exercised moral understanding, since that is the ‘aim’ of those states. So it can be perfectly rational to form moral judgements without trusting the testimony and judgements of others,

because making up your own mind is the only way to exercise your moral understanding. 'Epistemology' for moral uliefs with regard to moral testimony and related matters (including moral disagreement) is going to look rather different from the standard epistemology for (narrow) belief. The study of the relationship between our cognitive states and the world, epistemology in the broadest sense, is a bigger subject than we realized, of which philosophers have so far studied only one part, and not necessarily the most important part either.<sup>31</sup>

**(p.24)** References

Bibliography references:

Aristotle. 2002. *Nicomachean Ethics*, trans. S. Broadie and C. Rowe. Oxford: Oxford University Press.

Blackburn, S. 1998. *Ruling Passions*. Oxford: Clarendon Press.

Brink, D. O. 1989. *Moral Realism and the Foundations of Ethics*. New York: Cambridge University Press.

Coady, C. A. J. 1992. *Testimony*. Oxford: Oxford University Press.

Copp, D. 2001. 'Realist-Expressivism: A Neglected Option for Moral Realism', *Social Philosophy and Policy* 18: 1–43.

Davidson, D. 1980. 'Actions, Reasons and Causes', in his *Essays on Actions and Events*. Oxford: Oxford University Press, 3–19.

Driver, J. 2006. 'Autonomy and the Asymmetry Problem for Moral Expertise', *Philosophical Studies* 128: 619–44.

Engel, P. 2004. 'Truth and the Aim of Belief', in D. Gillies (ed.), *Laws and Models in Science*. London: King's College Publications, 77–97.

Gendler, T. 2008. 'Alief and Belief', *Journal of Philosophy* 105: 634–63.

Gibbard, A. 2003. *Thinking How to Live*. Cambridge, MA: Harvard University Press.

Glover, J. 1999. *Humanity: A Moral History of the Twentieth Century*. London: Cape.

Greene, J. D. 2007. 'The Secret Joke of Kant's Soul', in W. Sinnott-Armstrong (ed.), *Moral Psychology*, vol. 3: *The Neuroscience of Morality: Emotion, Disease, and Development*. Cambridge, MA: MIT Press, 35–79.

Grimm, S. 2006. 'Is Understanding a Species of Knowledge?' *British Journal for the Philosophy of Science* 57: 515–35.

Haidt, J. 2001. 'The Emotional Dog and its Rational Tail: A Social Intuitionist Approach to Moral Judgment', *Psychological Review* 108: 814–34.

Haidt, J. and Bjorklund, F. 2008. 'Social Intuitionists Answer Six Questions about Moral Psychology', in W. Sinnott-Armstrong (ed.), *Moral Psychology*, vol. 2: *The Cognitive Science of Morality: Intuition and Diversity*. Cambridge, MA: MIT Press, 181–217.

Harman, G. 2000. 'Moral Philosophy Meets Social Psychology: Virtue Ethics and the Fundamental Attribution Error', in his *Explaining Value and Other Essays in Moral Philosophy*. Oxford: Oxford University Press, 165–78.

Hills, A. E. 2009. 'Moral Testimony and Moral Epistemology', *Ethics* 120: 94–127.

Hills, A. E. 2010. *The Beloved Self: Morality and the Challenge from Egoism*. Oxford: Oxford University Press.

Hopkins, R. 2007. 'What is Wrong with Moral Testimony?', *Philosophy and Phenomenological Research* 74: 611–34.

Horgan, T. and Timmons, M. 2006. 'Cognitivist Expressivism', in Horgan and Timmons (eds.), *Metaethics after Moore*. Oxford: Oxford University Press, 255–98.

Hussar, K. M. and Harris, P. L. 2010. 'Children Who Choose Not to Eat Meat: A Study of Early Moral Decision-Making', *Social Development* 19: 627–41.

Jones, K. 1999. 'Second-Hand Moral Knowledge', *Journal of Philosophy* 96: 55–78.

Kamtekar, R. 2004. 'Situationism and Virtue Ethics on the Content of Our Character', *Ethics* 114: 458–91.

**(p.25)** Kant, I. 1991. *Metaphysic of Morals*, trans. M. Gregor. Cambridge: Cambridge University Press.

Kant, I. 1998. *Groundwork of a Metaphysic of Morals*, trans. M. Gregor. Cambridge: Cambridge University Press.

Kitcher, P. 2002. 'Scientific Knowledge', in P. Moser (ed.), *The Oxford Handbook of Epistemology*. Oxford: Oxford University Press, 385–407.

Kohlberg, L. 1973. 'The Claim to Moral Adequacy of a Highest Stage of Moral Judgment', *Journal of Philosophy* 70: 630–46.

- Kvanvig, J. 2003. *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.
- Lipton, P. 2004. *Inference to the Best Explanation*, 2nd edition. London: Routledge.
- McDowell, J. 1979. 'Virtue and Reason', *The Monist* 62: 331–50.
- McGrath, S. 2011. 'Skepticism about Moral Expertise as a Puzzle for Moral Realism', *Journal of Philosophy* 108: 111–37.
- Owens, D. 2003. 'Does Belief Have an Aim?', *Philosophical Studies* 115: 283–305.
- Pritchard, D. 2005. 'Knowledge, Understanding and Epistemic Value', in A. O'Hear (ed.), *Epistemology*. Cambridge: Cambridge University Press, 19–43.
- Ridge, M. 2006. 'Ecumenical Expressivism: Finessing Frege', *Ethics* 116: 302–36.
- Ridge, M. 2007. 'Ecumenical Expressivism: The Best of Both Worlds?', in R. Shafer-Landau (ed.), *Oxford Studies in Metaethics*, vol. 2. Oxford: Oxford University Press, 51–76.
- Sabini, J. and Silver, M. 2005. 'Lack of Character? Situationism Critiqued', *Ethics* 115: 535–62.
- Schroeder, M. 2008. *Being For: Evaluating the Semantic Program of Expressivism*. Oxford: Oxford University Press.
- Schroeder, M. 2009. 'Hybrid Expressivism: Virtues and Vices', *Ethics* 119: 257–309.
- Setiya, K. 2007. *Reasons without Rationalism*. Princeton, NJ: Princeton University Press.
- Sinclair, N. 2006. 'The Moral Belief Problem', *Ratio* 19: 249–60.
- Smith, M. 1987. 'The Humean Theory of Motivation', *Mind* 96: 36–61.
- Van Roojen, M. 2013. 'Moral Cognitivism vs. Non-cognitivism', *Stanford Encyclopedia of Philosophy* <<http://plato.stanford.edu/entries/moral-cognitivism/>>.
- Velleman, J. D. 2000. 'On the Aim of Belief', in his *The Possibility of Practical Reason*. Oxford: Oxford University Press, 244–82.
- Wedgwood, R. 2002. 'The Aim of Belief', *Philosophical Perspectives* 16: 267–97.

Williams, B. 1981. 'Internal and External Reasons', in his *Moral Luck*. Cambridge: Cambridge University Press, 101–13.

Williams, B. 1995. *Making Sense of Humanity*. Cambridge: Cambridge University Press.

Williamson, T. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.

Woodward, J. B. 2003. *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.

Zagzebski, L. 2001. 'Recovering Understanding', in M. Steup (ed.), *Knowledge, Truth and Duty: Essays on Epistemic Justification, Responsibility and Duty*. Oxford: Oxford University Press, 235–51.

Notes:

(<sup>1</sup>) Questions concerning the nature of moral judgement can relate to the proper interpretation of moral sentences (a semantic thesis) or the nature of the state of mind which is typically expressed by moral sentences (a psychological thesis). I will be discussing the latter in this chapter (see Van Roojen (2013) for further discussion), though the debates are related. Many authors mentioned in this chapter primarily discuss the former (though their views have implications for the latter question too).

(<sup>2</sup>) See Van Roojen (2013). Non-cognitivists include Blackburn (1998) and Gibbard (2003). Cognitivists include Brink (1989). Hybrid theorists include Copp (2001) and Ridge (2006, 2007); for discussion and assessment, see Schroeder (2009).

(<sup>3</sup>) There are of course religious and moral traditions that do endorse relying on moral testimony. I have argued at length elsewhere that these traditions are missing something of moral significance by accepting moral testimony (Hills 2009).

(<sup>4</sup>) See Jones (1999), Driver (2006), Hopkins (2007), Hills (2009), and McGrath (2011).

(<sup>5</sup>) Is this how we actually do make moral judgements? As far as I am aware, no psychologist has investigated directly our responses to moral testimony. The very well-known theories of Greene and Haidt, for instance, do not discuss pure moral testimony directly (e.g., in Greene (2007), Haidt (2001), and Haidt and Bjorkland (2008)). Kohlberg does say that some moral judgements may be made 'autonomously' suggesting that they are typically not based on pure testimony (Kohlberg (1973)). However, there is some suggestive evidence in Hussar and Harris (2010). This is a study of children under the age of 10 who are vegetarians though their families eat meat. When asked why they were

vegetarian, these children tended to cite animal welfare as the reason (rather than the taste of meat or health considerations) and would condemn themselves for eating meat as they would other moral transgressions. This suggests that their vegetarianism is based on a moral judgement that eating meat is morally wrong, and that they have ignored any moral testimony from their family that meat eating is morally acceptable (it is consistent with the study that they have based their view on moral testimony from non-family members, but also consistent with it that they have not and chosen to use their own moral judgement instead). Though this study is not conclusive, it does suggest that sometimes children may choose to make up their own minds about moral questions rather than trust moral testimony, even from a source like their parents whom they normally trust. Nevertheless, in this chapter I do not make assumptions about how moral judgements are actually formed, rather I am interested in how many people (including many moral philosophers) have thought that they should be formed.

(<sup>6</sup>) Williams (1995: 205). Kant's views about moral testimony are of course closely connected to his views about autonomy, expressed in all his ethical texts, for instance in the *Groundwork* at 4.425–7 and especially 4.431. He specifically discusses teaching ethics in the *Metaphysics of Morals* and rejects a lecturing model in which the pupil trusts her teacher in favour of a 'catechism' in which the teacher asks the pupil questions in order to draw out her understanding (6.478–84).

(<sup>7</sup>) For instance, *NE* 1144a6–8, *NE* 1113a32–3.

(<sup>8</sup>) I discuss this further in Hills (2009).

(<sup>9</sup>) I discuss whether there is ever good reason to treat the opinion of others as evidence for moral judgements in much more detail in Hills (2010), especially chapter 10.

(<sup>10</sup>) Note that this is true even for an internalist about knowledge or justification, who requires the justification for a belief to be accessible to the subject. If you have good evidence that someone is trustworthy (e.g., she has a good track record), then even an internalist accepts that you are justified in accepting her pure testimony, and by doing so, you can gain knowledge.

(<sup>11</sup>) With regard to autonomy, it is notoriously unclear what autonomy requires, and whether it is essential to morality. One possibility is that you can only be bound by judgements that you make for yourself, and you can only have moral obligations if you are bound by them. This would imply that at least one class of moral judgements (those about obligations) need to be self-given. But the argument is not convincing. Why cannot recognizing an obligation on the basis of someone else's testimony bind you too? With regard to motivation, it is certainly not obvious that we could not be motivated by a 'second hand' moral

judgement. And if motivation did come with the judgement, there seems to be no reason not to take it on trust.

(<sup>12</sup>) Davidson appears to contrast believing with other attitudes: ‘knowing, perceiving, noticing, remembering’ (Davidson 1980: 3). But since it is natural to think that knowing, perceiving, etc. all involve belief, it is natural to assume that the parenthesis distinguishes different kinds of belief or knowledge (e.g., from different sources) rather than different cognitive states that are to be distinguished from belief and knowledge.

(<sup>13</sup>) In response to this variety, many different non-cognitivist views have been set out and defended. To take just a sample, it has been claimed that a moral judgement that x is wrong is a combination of attitudes of disapproval of x and disapproval of those who fail to share this disapproval (Blackburn 1998); a positive attitude in favour of blaming for the action in question (Schroeder 2008); and a planning state concerning feelings of guilt and resentment (Gibbard 2003). These theories all concern desires in the broad sense but not desires in the narrow sense.

(<sup>14</sup>) The recognition of other kinds of cognitive state (e.g., in Velleman 2000: 244–82) has not much influenced characterizations of cognitivism. There has been some discussion about the possibility of two kinds of moral belief, a ‘minimal’ moral belief which even expressivists can take moral statements to express, and a more robust kind of moral belief, which they cannot (Blackburn 1998; Sinclair 2006; see also Horgan and Timmons 2006). Minimal belief may ‘aim’ at truth in a minimal sense, or (in the case of Horgan and Timmons) may not be a descriptive state at all. Nevertheless, all of these discussions assume that the cognitive mental states in question are beliefs, and so do not address the objection to cognitivism in which I am interested (in addition they face further issues, as to whether minimal beliefs or non-descriptive beliefs can capture all the important features of moral judgements).

(<sup>15</sup>) There is a large literature attempting to explicate the sense, if any, in which beliefs aim at the truth. See for example Velleman (2000), Wedgwood (2002), and Owens (2003). Owens suggests that beliefs may aim at knowledge rather than truth, as does Williamson (2000) and Engel (2004).

(<sup>16</sup>) I aim to be as neutral as possible in characterizing knowledge, belief and evidence here. For instance, this thesis has no implications for the question of whether knowledge or belief is a ‘conceptually prior’ notion. I am also neutral on the question of what evidence is (except that I claim that testimony that p can be evidence that p) and I am not endorsing evidentialism (i.e., I do not assume that your belief is justified iff it fits your evidence).

(<sup>17</sup>) This account of moral understanding is based (with a few minor revisions and clarifications) on that in Hills (2009, 2010) where it is discussed in much more depth.

(<sup>18</sup>) There is a widespread though not universal view that understanding why p is the same as knowing why p (Kitcher (2002), Woodward (2003: 179), Lipton (2004: 30), and Grimm (2006) agree, Zagzebski (2001), Kvanvig (2003), and Pritchard (2005) do not). My view of understanding is not exactly the same as any of these (see Hills (2009, 2010) for more discussion).

(<sup>19</sup>) A few more brief remarks about explanation. First, explanations can be more or less full. To understand why m, you do not need to have the fullest possible explanation, an explanation 'all the way down'. How many of the morally relevant features you need to be sensitive to, how full an explanation you need to be able to give may depend on your circumstances. Explanations of moral truths that involve moral claims (e.g., it is wrong to raise sales tax because doing so unfairly burdens the worse off) are perfectly legitimate. It is possible that ultimately all moral claims may be explicable in non-moral terms (because they ultimately are grounded on natural features of the situation) but I do not suppose that in order to understand why m you must be able to give an explanation of it in non-moral terms.

(<sup>20</sup>) In addition, if you judged that p on the basis of q in the face of widespread disagreement, you might understand why p without knowing that p. For instance, suppose that many people whose judgement you had previously thought perfectly good, said that not-p, this testimony would be a defeater for knowledge—it would defeat whatever grounds you had for knowing that p. So even if you continue to believe that p, you could not know that p (or know why p). But provided that your judgement that p (because q) was true, you could have understanding, even if that judgement was made in the face of misleading evidence. I discuss this sort of case at great length in Hills (2010).

(<sup>21</sup>) It is sometimes said that a virtuous person may 'just see' what it is right to do. And when you 'just see' that an action is right, you do not deliberate about whether it is right beforehand (and you may not even be aware exactly why it is right: McDowell (1979: 332)). But even if that is right, it just shows that you do not need to exercise all of the abilities essential to moral understanding any time that you decide what to do. You still need to make your judgement for the right reasons, i.e., to have and to use abilities (iii) and (iv), and you might well need the others as well, to explain and to justify your action to others, or to give advice.

(<sup>22</sup>) This is why moral uliefs are said to aim at 'exercised' moral understanding, rather than merely that they aim at moral understanding. As I mentioned earlier, moral understanding consists in making a judgement (that p because q (or that q

is why p)) and in having certain abilities. You might have these abilities, but nevertheless make your moral judgements on other grounds: testimony for example. But in order to be virtuous, it is important not just to have these abilities, but to use (some of) them to make your moral judgements. And if you exercise moral understanding in making your moral judgements—or even if you try to—you are forming moral *uliefs*, not beliefs.

(<sup>23</sup>) What if it was not a case of pure moral testimony and the priest explained his reasons? This is a more complicated example that I do not have space to consider here, but see Hills (2009) for further discussion.

(<sup>24</sup>) See Jones (1999) for a good example. This is similar to the familiar problem of when it is right and when it is not for a non-virtuous person to try to act as the virtuous person does. The best-known examples concern non-virtuous non-cognitive attitudes. For instance, you might be more prone to anger or jealousy than a virtuous person, and it would be right for you to avoid certain situations, which the virtuous person has no need to avoid. Similarly, sometimes it might be right for you not to make moral judgements in the way a virtuous person does, because doing so would be likely to end in failure.

(<sup>25</sup>) Perhaps an ‘*alief*’ as described in Gendler (2008) might be a cognitive state (though it also has non-cognitive and affective aspects too), and more familiarly, suppositions and hypotheses might qualify as well (see Velleman (2000)).

(<sup>26</sup>) Both arguments are made at greater length in Hills (2009, 2010).

(<sup>27</sup>) Other interpretations are advanced by Harman (2000) and Kamtekar (2004), amongst others.

(<sup>28</sup>) Of course there is a complicating factor in the Milgram experiment, namely that the shocks are not real and so no one at all is suffering. But anyone who was aware of this would not be an appropriate subject for the experiment. And a virtuous person in a real-life situation of the type the experiment is trying to model would not agree to give significant shocks to the victim. Glover (1999) recounts similar real-life examples, including the My Lai massacre, in which many ordinary US soldiers obeyed an order to kill unarmed civilians and only a few refused. According to Glover (1999: 333) one of the whistleblowers on the massacre, Ronald Ridenhour, had been the subject of a repeat of the Milgram study at Princeton, and refused to give any shocks at all.

(<sup>29</sup>) It does not follow from this that a morally virtuous person never defers to anyone else’s moral judgement. There may be specific circumstances in which she has good reason to think that her own judgement is impaired and she should trust someone else, but in these circumstances she is not acting ‘characteristically’. Note also that even if a morally virtuous person would not defer to someone else in a particular situation, it does not follow that the rest of

us should not either. We should not always try to act as the virtuous person does—our own shortcomings can make such attempts disastrous. Making moral judgements as a virtuous person does is always a moral ideal, and we should do so where possible, but if our judgement is very bad (or the situation very unfavourable) we should trust testimony, rather than trying to exercise moral understanding.

(<sup>30</sup>) There is a well-known argument that no state can have both the direction of fit of a belief and that of a desire (Smith 1987: 54f.). I am not convinced that this argument succeeds, for reasons similar to those given by Setiya, in the course of his argument that intentions are mental states with both belief-like and desire-like components (Setiya 2007: 48–51). Moreover, this chapter has shown that mental states can be ‘belief-like’ in some respects but not in others. Perhaps there could be a mental state that is belief-like in some respects—sufficiently belief-like to qualify as a cognitive state—but also desire-like in some respects. But I cannot pursue this interesting possibility any further here.

(<sup>31</sup>) Thanks to audiences at Oxford, BSET, SPAWN Stirling, and the Madison Metaethics conference where this essay was first presented and especially Rachel Cohon and Debbie Roberts.

Access brought to you by: