

Moral Testimony and Moral Epistemology*

Abstract

According to *pessimists* about moral testimony, there are strong reasons not to form moral beliefs on the say-so of others. I defend pessimism by setting out a conception of what I call *moral understanding*. I argue that moral understanding is different from moral knowledge and that it is an essential part of good character and of morally worthy action. I argue that if we are aiming to acquire and to use our moral understanding, as we should, it is best neither to trust moral testimony nor to defer to moral experts but that taking moral advice can be worthwhile.

1. A puzzle about moral testimony.

The status of *moral testimony* is controversial.¹ Consider the following example:

Eleanor has always enjoyed eating meat, but has recently realized that it raises some moral issues. Rather than thinking further about these, however, she asks a friend who tells her that eating meat is wrong. Eleanor knows that her friend is normally trustworthy and reliable, so she believes her and accepts that eating meat is wrong.

Many people believe that there are strong reasons not to form moral beliefs on the say-so of others, as Eleanor does. I will call these people *pessimists* about moral testimony. Pessimists do not think that moral testimony is entirely worthless. They agree that it is acceptable for children to learn most of their moral beliefs through testimony. And adults may reasonably base some beliefs on testimony, for example, about relevant non-moral issues (who did what to whom) from which they can form their own moral views. Asking for other people's advice about moral matters and taking that advice seriously is clearly legitimate. But once you have reached maturity as an adult and have the ability to think about moral questions by yourself – as Eleanor can – it seems, at least to pessimists, that you have strong reasons to do so, indeed that refusing to do so is unacceptable.

On the other hand, *optimists* think that trusting moral testimony is just like trusting testimony about an ordinary non-moral matter of fact. To defend this claim, they often appeal to moral epistemology. By trusting testimony, you can gain knowledge, both moral and non-moral. And if you can gain moral knowledge by trusting moral testimony, what reasons, let alone strong reasons, could there be against doing so? Of course, no one thinks that you should *always* trust moral testimony: you certainly should not do so if your interlocutor is obviously unreliable. But whenever the speaker is trustworthy, experienced and

knowledgeable, trusting what they say is perfectly legitimate; according to optimists, Eleanor has every reason to trust moral testimony, and her doing so is entirely acceptable.²

There is a very similar dispute about moral expertise. Everyone agrees that there are *experts* regarding ordinary non-moral matters of fact, people who are more experienced and more knowledgeable than the rest of us. We often *defer* to these experts: we form our beliefs on the basis of their judgements alone, without making our own assessment of the reasons for their judgements (perhaps without even being aware of those reasons). And we often have very good reasons to defer: we thereby acquire knowledge. According to some ethical traditions, there are experts regarding morality, just like any other subject, and the rest of us should defer to them. The experts may also be religious authorities, for example, priests or rabbis. These traditions value submission and obedience to the experts, rather than making up your own mind about morality.

By contrast, many people think that there are no moral experts, and perhaps that there could be none. Of course, some people have thought for a long time about particular moral issues and are skilful at formulating moral arguments; some of these people are moral philosophers. People frequently ask for advice about moral questions, and they might prefer to consult people who had thought about the issues carefully. But many adults do not defer to the opinions of others in moral matters as they do in ordinary factual ones and they think that they rarely, if ever, have reason to do so. Bernard Williams, for example, is briskly dismissive of the very idea of deferring to a moral expert:

There are, notoriously, no ethical experts.... Anyone who is tempted to take up the idea of there being a theoretical science of ethics should be discouraged by reflecting on what would be involved in taking seriously the idea that there were experts in it. It would imply, for instance, that a student who had not followed the professor's reasoning but had understood his moral conclusion might have some reason, on the strength of his professorial authority, to accept it... These Platonic implications are presumably not accepted by anyone.³

A possibility so absurd, he claims, that only Plato believed it. According to Williams, this is a striking difference between the epistemology of morality and the epistemology of ordinary facts. There are no moral experts, or at least none to whom we should defer; but plenty of experts in other fields.

There are some obvious and uncontroversial reasons why we might not defer to moral experts or trust what they say. First, it may be more difficult to identify experts in ethics rather than other subjects. If we tried to establish the qualities essential to a moral expert, they would prove extremely contentious. Of course, a moral expert must be sensitive to morally

important features, but what kind of sensitivity, to which features, is needed? The kind of people whom we might regard as moral experts – those who have spent much of their lives thinking about moral questions, for example – often disagree greatly amongst themselves about the right answers. Furthermore, moral experts, even if they have moral knowledge, may not transmit it reliably to others. And the stakes are often very high in situations where we face moral questions. When your decision is so important, you may well prefer to trust your own judgement rather than defer to someone who may or may not be an expert.⁴

But although it may be difficult to identify moral experts, this is not a particularly deep or in principle an insurmountable problem. Few people, if any, can properly be described as experts on all moral matters. But to discover someone who has better judgement or more experience than you about some single moral issue and who seems generally trustworthy is not that difficult. Why not take her word on at least that moral question?⁵

Pessimists about moral testimony believe that the problem is much more serious. Williams does not think it is simply difficult to tell if your ethics professor is an expert or that she is trustworthy. There is no reason to defer to her, to put your trust in her testimony, he seems to imply, even if you know she is very likely to be right. But then why should it be acceptable to take moral advice from her, as it plainly is? The puzzle about moral testimony, then, is why you have reason neither to defer to moral experts nor to trust moral testimony, if indeed you do, given that taking moral advice is obviously acceptable and may even sometimes be required.

I am going to defend pessimism about moral testimony, that is, I am going to argue that there are circumstances in which you have no reason to trust moral testimony (in fact, you have reason *not* to put your trust in it), even if your interlocutor is reliable and trustworthy regarding the matter in question and you know her to be so. But I will begin by agreeing with the optimists: I will simply accept that you can acquire moral knowledge through testimony.⁶ This claim has been defended recently and I have nothing to add to those arguments.⁷ But if optimists are right that trusting moral testimony can give us moral knowledge, their question is pressing: what reasons could there be not to trust moral testimony or defer to moral experts? I will answer this question by taking a new look at moral epistemology. I will argue that 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.

I set out my conception of moral understanding and explain how it differs from moral knowledge (both knowledge that p, and knowledge why p) in section 2. In section 3, I argue that moral understanding is extremely important: it plays a vital role in good character and in morally worthy action. I then show (in section 4) that, given the importance of our acquiring and using moral understanding, we have strong reasons neither to trust moral testimony nor to

defer to moral experts, though taking moral advice is both acceptable and often very useful.⁸ Ethical traditions which defend trusting moral testimony and deferring to moral experts are, I will argue, missing something of vital moral importance.

2. What is moral understanding?

Consider some explicitly moral propositions:

It is morally right to help those in need.

I have moral reason not to lie.

Action X is morally wrong.

What is it to *understand why* any of these are true?⁹ Contrast the following cases:

Sarah is five years old. Her parents have told her not to lie, that lying is naughty and wrong. She accepts what they say.

Peter is fifteen. He has heard many people say that helping those in need is morally right, but he does not agree. He thinks that the strong are entitled to trample on the weak and should not be distracted by talk of needs.

Mary believes that she has moral reason not to lie because lying to others fails to respect them and in the long-run tends to make them unhappy. She can see that there are differences between lying and not telling the whole truth. She thinks that lying for your own benefit is normally wrong, but it is harder to say whether lying to someone in a way that makes her happy is ever acceptable, and this may depend on the exact circumstances of the case. She can apply these beliefs to new situations and judge which action is morally right.

Let us suppose that Mary's beliefs, sketched above, are correct. She understands why she has moral reason not to lie. Peter by contrast has a false belief – that he has no moral reason to help those in need – so he cannot have moral understanding. In the sense in which I am interested here, moral understanding is factive: you cannot understand why p if p is false.¹⁰

Moral understanding, in this sense, is not transparent. You might *think* that you understand perfectly well that refusing to help the needy is morally acceptable, but of course

you do not: you cannot. It is fairly easy to see how you might fail to understand something that you thought you understood: Peter thought he understood why it was not right to help the needy but he was wrong. It is perhaps less clear that you can have more moral understanding than you think. But this is possible. Suppose that you are faced with a very difficult moral situation. You form a judgement about what to do based on what you take to be relevant. You are not at all sure that you have it right and are concerned that you may have missed out some relevant factors or classed some of them wrongly. In fact, you have correctly identified all the factors and weighed them appropriately. It is possible that you understand why this action is right and alternatives are wrong, even though you do not realize that you do.

In order to have moral understanding in this sense, you do not have to deliberate extensively about each case and you may not need to deliberate at all. Peter may have thought about what is morally right and wrong much more than Mary has, but she has moral understanding and he does not. In more complex situations, careful and lengthy deliberation may be necessary.

Peter clearly cannot understand why helping the needy is not morally right, for he is mistaken about that. What about Sarah? She correctly believes that lying is morally wrong. But she does not understand why. In general, young children tend to lack moral understanding, even if they have been brought up with correct moral beliefs. Sarah, like other young children, is relying on her parents' moral opinions. Unlike Mary, she has not yet developed her own judgement.

Contrast Mary, too with Eleanor. Suppose that Eleanor trusts her reliable friends, and thus knows that eating animals is wrong. But whereas Mary was able to offer an explanation of why lying is wrong, Eleanor is not at all sure why eating meat is wrong (is it eating or killing that is wrong? Do animals have rights or is it a question of maximizing welfare?) In short, Eleanor may *know that* eating animals is wrong, but she does not *understand why*.

Understanding why p, therefore, is different from knowing that p. According to some accounts of understanding, however, there is a close connection between understanding and a different type of knowledge: understanding why p is the same as knowing why p.¹¹

Suppose that Eleanor's friend offered her an explanation: eating animals is wrong because of the suffering of animals under modern farming methods. If this is correct 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.¹² This suggests that understanding why p

differs both from knowledge that p and knowledge why p.¹³ It is possible to have moral knowledge without moral understanding. But what exactly is the difference?

If you understand why X is morally right or wrong, you must have some appreciation of the reasons why it is wrong.¹⁴ Appreciating the reasons why it is wrong is not the same as simply *believing that* they are the reasons why it is wrong, or even *knowing* that they are the reasons they are wrong. Moral understanding involves a *grasp* of the relation between a moral proposition and the reasons why it is true.¹⁵

You cannot really understand why p is true if p and the reasons why p are the only things about the subject of which you are aware: it is not possible to understand why some isolated fact is true. If you have this kind of appreciation of moral reasons you must have, at least to some extent, a systematic grasp of morality. This is not, of course, to say that you need to have a grasp of anything that could possibly be morally important. But if you understand why you have moral reason to help the needy, for instance, you must have some awareness that the needs of other people are morally important and some grasp of their relative importance compared to other moral and non-moral considerations (such as self-interest).¹⁶ You must also be able to make appropriate judgements in similar cases, for example, that X' is morally wrong for the same reason, but Y isn't. One sign that Eleanor and Sarah have no moral understanding is precisely that they are unable to draw conclusions about similar cases. If you do have the relevant abilities, you can generate new true moral beliefs (and perhaps new moral knowledge) without having to put your trust in anyone else's judgement. Your ability to draw the right conclusion or give the right explanation in new cases is not a matter of luck: it is explained by your appreciation of the reasons why p. Since you cannot understand why p unless p is true and you have the appropriate grasp of the reasons why p, understanding why p must be something that can be credited to you: understanding therefore is always a kind of achievement for you, though of course others can help you on the way, and sometimes their help and advice will be invaluable.¹⁷

The grasp of the reasons why p that is essential to understanding involves a number of abilities: to understand why p, you need to be able to *treat* q as the reason why p, not merely *believe* or *know that* q is the reason why p.¹⁸ If you understand why p (and q is why p), then in the right sort of circumstances you can successfully:

- (i) follow an explanation of why p given by someone else
- (ii) explain why p in your own words
- (iii) draw the conclusion that p (or that probably p) from the information that q
- (iv) draw the conclusion that p' (or that probably p') from the information that q' (where p' and q' are similar to but not identical to p and q)
- (v) given the information that p, give the right explanation, q;
- (vi) given the information that p', give the right explanation, q'.¹⁹

To understand why p, you have to have the abilities (i)-(vi) to at least some extent. Of course you can have these abilities to a greater or lesser degree, which may make it tempting to say that moral understanding comes in degrees. You have minimal moral understanding if you correctly believe that q is why p and you can follow an explanation of why p. You have greater understanding the more you fulfil (i)-(vi), and you have full understanding if you have (i)-(vi) to the greatest extent.²⁰ Alternatively, there might be a cut-off point before which you do not count as having moral understanding, after which you do. This cut-off point might be contextually determined, i.e. in one context you might need (i)-(vi) to only a minimal degree to count as having moral understanding, a different context might have higher standards, so that you needed to have (i)-(vi) to a greater extent to have moral understanding. I will not settle this question here, as the differences between the two conceptions of moral understanding will not much matter.

You can know why p without having these abilities – that is one reason why understanding why p is not the same as knowing why p. But it may be tempting to think that having the abilities is simply to have extra pieces of knowledge. For example, perhaps to understand why p, you need to know why p (where q is why p) *and* to know if q' rather than q were true, then p' rather than p would be true. Moral understanding would not be identical to knowing why p, but it would be the same as a combination of knowing why p plus some other pieces of propositional knowledge.

But I think that having these abilities is not the same as having extra pieces of knowledge. Gaining this extra knowledge may help you acquire the requisite abilities, but you might have the extra pieces of knowledge without having the kind of good judgement that enables you to generate new true moral beliefs yourself. Surely no extra piece or pieces of knowledge guarantee that you have these abilities.

A second reason for doubting that understanding why p is some combinations of knowledge why p plus some extra pieces of propositional knowledge is that moral understanding has a different relationship to epistemic luck than knowledge does.

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

the ability to draw the conclusion that he was evil from the reasons why he was evil and to do the same in similar cases. So it seems that you can have moral understanding why p without having knowledge why p.²¹

If this is right, moral understanding is distinct from propositional knowledge, but it has, I think interesting similarities with knowing how: they both have close connections to certain sorts of ability. Of course the relationship between propositional knowledge and knowing how is itself a matter of dispute. Whilst according to some accounts, knowing how and propositional knowledge are distinct, there are well-known arguments that knowing how is a species of propositional knowledge.²² If these arguments apply to understanding why p as well as to knowledge how, one might conclude that moral understanding must be a species of propositional knowledge too.²³

I am inclined to think that understanding why p is not a kind of propositional knowledge, for the two reasons that I have given here. But I do not intend to press this, since I do not think that ultimately it will matter whether moral understanding is classed as propositional knowledge or not. Suppose that Stanley and Williamson, for example, are right that knowing how is a species of propositional knowledge and their arguments also apply to understanding why p (as I conceive of it here). Nevertheless, they concede that there are some differences between the two. They make a distinction between knowing a proposition “under a practical mode of presentation” – as you do when you have know how - and knowing it under a different mode (say, a “theoretical mode”). They say that it is difficult to characterize what it is to know a proposition “under a practical mode of presentation” but it undoubtedly involves the possession of certain complex dispositions. Given this distinction, it seems that you might know propositions under one mode of presentation, but not the other (because you know it without having the relevant dispositions and so lack know how; or because you have the relevant dispositions and so have know how but do not have knowledge in another form). In which case moral understanding may include knowing propositions under a practical mode of presentation, but not necessarily under a theoretical mode of presentation – so moral understanding would be a form of propositional knowledge but very unlike ordinary propositional knowledge. As we shall see, the points that I want to make about moral understanding follow whether or not understanding why p is a form of propositional knowledge, provided it is propositional knowledge of a sufficiently special kind.

I will have more to say about the differences between moral knowledge and moral understanding in section 4, with regard to moral testimony, but first I want to explain why moral understanding is important.

3. The value of moral understanding.

Why should we seek moral understanding? I will argue that moral understanding is valuable in a number of ways. In the first place, it can be useful because it is a route, perhaps the only route, to reliably doing right.

A. Reliably doing right.

Moral understanding might be important because you yourself having some systematic grasp of moral reasons might in practice be the only way that you could reliably do the right thing. Whilst in principle, your moral instincts might be infallible so that you always instinctively chose the morally right action or you might keep your moral guru by your side at all times to advise you, in practice neither of these is likely to happen. Moral decisions are often complicated. Moral reasons can be difficult to assess and interact in quite complex ways. Small differences in factual situations can make significant moral differences. You often have to make judgements of what to do quite quickly, so you do not have the time to find and consult with an expert. There is simply no one to ask; you have to make the decision on your own and you will make a good decision only if you have moral understanding – the ability to make accurate judgements in new circumstances - or luck.

So moral understanding is worth having as a means to right action. Moreover, there is a second way in which moral understanding can be useful: in order to justify your actions to others.

B. Justifying yourself to others.

Suppose you have been injured in a car accident and though you are my friend, I don't go to help you. You demand to know why. I reply that there are other people who are hurt more badly than you. You protest that you are my friend and they are not. I reply that the injuries of the others are potentially fatal whereas yours is not, and I must look after them first.

A core ethical practice is the exchange of reasons. According to Scanlon's contractualist theory, to be motivated morally precisely is to be motivated to find a justification for what we do.²⁴ Whilst I do not intend to defend contractualism generally, nor do I want to suggest that this is the only (or even the most important) kind of moral motivation, nevertheless, the practice of exchanging reasons and the motivation to find a justification that could not be reasonably rejected by others is clearly morally very important.

The requirement to be able to justify yourself involves at least being able to say what you take yourself to be doing and why you thought doing so was a good idea. You might say: I had to help those people who were strangers to me because they were very badly injured, whereas my friend's injuries were only minor. Giving a justification involves giving the

reasons why what you did was right. If you do not understand why your action was right, you are in a very awkward position.

Suppose your reliable friend has told you to not to cheat your customers because doing so is unfair. You believe her, but on your own behalf, you cannot really see anything wrong with enriching your shareholders at your customers' expense. There is a perfectly good sense in which you know why cheating your customers is wrong: you know that it is unfair. And yet you yourself have not grasped the connection between the wrongness of the action and the reasons why it is wrong. So you say to your customers what you were told – citing the unfairness of giving the wrong change as a justification for your action – but you cannot give an explanation in your own words, and you cannot reassure customers that under slightly different circumstances, you would treat them well (since you yourself are not able to work out what you would have moral reason to do in those circumstances). Without moral understanding, your ability to participate in the exchange of reasons is necessarily limited.

So moral understanding is important in part because being in a position to justify yourself to others is morally important. On the other hand, there may be limits to how well you can justify yourself to others even if you understand why what you did was right. One of the abilities associated with moral understanding is the ability to give your own explanations, but you need not be exceptionally articulate or persuasive in order to have moral understanding. In any case, even if the exchange of reasons is genuinely important, there may be a more basic reason why moral understanding is morally significant.²⁵ I will argue that moral understanding is essential to good character and to morally worthy action, that is, to right actions performed for the right reasons. The argument will turn on some subtle but important aspects of what it is to respond to moral considerations in forming your moral beliefs and in action. I will begin with the role of moral understanding in having a good character.

C. Virtue

What is it to have a good character, to have the virtues? Someone who is honest tends to tell the truth, to disapprove of lying and of breaking promises. Someone who is courageous is willing to stand up to danger when there is good reason to do so. In general, someone with a good character is disposed to feel, choose and act rightly.

Being virtuous is obviously a complex matter consisting of number of related dispositions. I want to separate out two aspects of good character: one cognitive and one non-cognitive. A virtuous person has *good motivations*: she wants to help others, to tell the truth and so on. These motivations need not be explicitly moral but sometimes they may be: she may want to do what is just, for example, under that very description. Good motivation is essential to having a good character, but it is obviously not sufficient. To be virtuous, you

have to care about helping others and telling the truth, but you cannot be fully virtuous if you consistently mistake what is just, honest, or kind in the circumstances. *Good judgement* is a second crucial component of virtue.²⁶

If you do not have good judgement, it is likely that you will actually act wrongly quite often, however worthy your motivations. And if you do the right thing, it will be an accident, because by chance you happen to have hit on the right action. Similarly, someone with the wrong motivations – someone who is essentially selfish, for example – will often act wrongly, and if she does the right thing, it will be by accident. Of course, you can also have bad judgement *and* bad motivations, and the two failures can be connected. For example, a very strong concern for your own interests can lead you to fail to recognize what it is morally right to do. But I will be focussing here on cases in which someone's bad judgement is not based on bad motivations.

Someone might suggest that good motivations and good judgement are essential to virtue because with them you will *reliably* act rightly and without them you will not. But I do not think that this is the most fundamental explanation for their importance. After all, it is not impossible for you to be in circumstances in which selfishness will lead you reliably to do the right action. You might even know that you are in those circumstances, and so know that you will reliably do right. But you are not thereby virtuous. The real problem is that if you care only for yourself, you are not really responding to moral reasons – even if you happen to be in circumstances in which reasons of self-interest are reliably aligned with moral reasons. A virtuous person's motivations must be *responsive to moral reasons*.

I think that the same is true for the cognitive aspects of virtue.²⁷ There are ways in which you might form beliefs about what is morally required of you, what would be just or kind in your circumstances that would reliably lead you to act rightly – and you might know that they would reliably lead you to act rightly – but you would still not be virtuous.²⁸ The most obvious is: you might ask someone else what to do. In the right circumstances, where you ask someone who is themselves virtuous and trustworthy, you will form true beliefs – indeed you will have knowledge – and if you have the right motivations, you will reliably do the right thing. But you will not be virtuous. Consider the following case:

The Incompetent Judge

Claire has just been appointed as a judge and is very anxious to sentence people justly. But she finds it exceptionally difficult to work out the just punishment for various offences, though she listens to the evidence presented carefully and tries her best to get the right answer. Luckily she has a mentor, a more experienced judge Judith who has excellent

judgement. Claire always consults with Judith and gives her decision in accordance with Judith's guidelines, offering Judith's explanation of why the sentence is just to the defendants.

Claire is not a good judge – she lacks important features of the virtue of judges, that is, of justice. This is not because she does what is just only occasionally and unreliably, or that it is only an accident when she acts rightly: she reliably gives the right sentence thanks to her good motivations and her knowledge. The problem is that she depends on someone else to tell her what to do: she is forming her beliefs in response to Judith, not in response to what is just.

Perhaps what is missing is extra knowledge: perhaps what she needs is knowledge *why* the sentence is just. But suppose that Judith tells her that 5 years in jail is the just sentence because of the type of offence committed. Claire then knows why the sentence is just. But she herself lacks the capacity to derive the conclusion that that sentence is just from the reasons why it is just in this and similar cases. So she is simply unable to base her belief that a 5-year term in prison is just on the reasons why it is just, she is still wholly relying on Judith's testimony. What Claire lacks, and what she needs in order to be just is *moral understanding*: she needs the *capacity* to base her beliefs about what is just on the reasons why it is just, and she needs to exercise that capacity: to form her beliefs about what is just *on the basis* of the reasons why it is just. Only then will she actually be forming her moral beliefs in response to moral reasons.

Suppose that action A is right because of reason R. If you have moral understanding, then you can derive the conclusion (that A is right) from the reasons why it is true (that R), and if you use your moral understanding, you do derive the conclusion from the reasons why it is true, hence the structure of your moral beliefs, the way in which you base one belief on another, mirrors the structure of morality.

Claire, by contrast, bases her belief that action A is right on the testimony of Judith. It follows that Claire's belief is not based on the reasons why it is true. Moreover, because she does not have moral understanding, she cannot derive the conclusion (that A is right) from the reasons why it is right, so even if she wanted them to, her moral beliefs could not mirror morality.

Might Claire reply that once Judith tells her why A is right, she can base her belief that it is just on the reasons why it is right, so the structure of her moral beliefs can mirror morality after all? But I think that this is not sufficient for her to claim that she is correctly oriented. I said earlier that if you understand why p, you can *treat* q as the reason why p; Claire lacks the abilities characteristic of moral understanding, and she cannot treat R as the reason why A is right (though she believes that it is the reason). Consider the basis of Claire's belief that A is right. Since she has no moral understanding, Claire is not able to derive the

conclusion that A is right from R without Judith telling her what to derive from what and helping her every step of the way, as shown by her inability to draw the conclusion in similar cases. Judith's testimony is the basis *both* of her belief that A is right *and* of her belief that R is why A is right. Claire's belief that A is right is explained by Judith's testimony and *there is no independent explanation for her belief*, as there would be if she had and used her moral understanding.²⁹ She is responding to testimony, not to moral reasons.

Why is responsiveness to moral reasons so important? Why isn't reliably doing right sufficient? One answer is that a mark of a virtuous person is that she is an *authority* into what is right.³⁰ Claire is not an authority on what is just: relying on the judgement of someone else – however dependable they are – is not a way for you to be an authority. It is, as it were, the difference between *being* a moral compass, and *having access to* a moral compass.

The second answer expands on this idea. Being a good person is not just about what you do. Reliably acting rightly is a part of having a good character, of course, but a good, virtuous person is someone whose *whole self* – her thoughts, decisions, feelings and emotions as well as her actions – is structured by her sensitivity to morality. This does not mean, of course, that a virtuous person is so consumed by morality that she thinks of nothing else. But it does mean that she is responsive to moral considerations in all aspects of her character, whenever they are relevant.

I will call this responsiveness to moral reasons, orientation. Someone who is appropriately oriented is sensitive to the features of actions that determine whether those actions are right or not, and to whether she has reason (and what strength of reason) to perform those actions. You may be sensitive to those features of actions without conceiving of them in moral terms. But when you form moral beliefs, appropriate orientation involves using your moral understanding. Knowledge from someone's testimony is inadequate, because you are relying on *their* being appropriately oriented towards moral reasons, without being so yourself. Without moral understanding, you cannot be properly oriented with regard to your moral beliefs: moral understanding is the cognitive, intellectual aspect of correct orientation to moral reasons, and as such it is an essential part of good character.

I have drawn an analogy between two ways that you can fall short of virtue: cognitive and non-cognitive. But the analogy may seem like a mistake, and a very deep one. Surely a selfish person and the incompetent judge are actually very different, and in a morally significant way. The selfish person is blameworthy in a way that Claire is not, if she really is well-motivated and is doing her best.

It is true that having poor judgement is not blameworthy in exactly the same way as failing to have the right motivations, though we might well blame Claire if, having failed to get in touch with Judith, she is forced to make her own judgement and she makes a decision that is very badly mistaken. Nevertheless, it is certainly true that we would for the most part

treat the selfish and the incompetent differently. But I do not think that this undermines my central claim that virtue requires that both your motivation and your judgement are responsive to moral reasons. We might not always blame the incompetent judge, but we would – and should – regard her as *less than fully admirable*. This is a negative evaluation of a different type than blame but it is still a negative evaluation, reflecting that lacking judgement is a different way of falling short of virtue than lacking right motivation, but that it is still a way of falling short.

D. Morally Worthy Action

So far I have argued that moral understanding is essential to good character. But I think that it also has a role to play in important kinds of moral action. There is a well-known distinction between doing the right action, and acting well or performing morally worthy actions. Your action is morally worthy only if it is a right action performed for the right reasons and I will argue that moral understanding is crucial to certain kinds of morally worthy action.

There is a connection between having a good character and performing morally worthy action: virtue can reasonably be regarded as a disposition to perform morally worthy actions. But you can perform morally worthy actions without being virtuous, because you do the right thing for the right reasons, though you do not habitually do so and you are not disposed to do so in other circumstances. And you can have the virtues without performing a morally worthy action on an occasion when you uncharacteristically fail to respond to moral reasons.

The best-known examples illustrating the difference between right action and morally worthy action involve different types of motivation. Kant famously drew a distinction between two shopkeepers. One gives the right change to his customers only for the sake of his reputation. He is not really acting well: he is doing what is morally right, but only because it is in his interests to do so. Whereas the shopkeeper who gives the right change because doing so is fair, in order to treat his customers with respect, is acting well and his action has *moral worth*.³¹ But I will argue that your beliefs and the ground of those beliefs are also crucial to acting rightly for the right reasons. Just as virtue had motivational and cognitive components, morally worthy action does too.³²

It is widely accepted that in order to act well, it is not sufficient that you do the right thing, because you (correctly) think it is morally right and you want to do the right thing, as an example from Nomy Arpaly illustrates:

The Extremist

Ron is an extremist, believing that killing a person is not generally immoral but that killing a fellow Jew is a grave sin. Ron would like to kill Tamara, but he refrains from doing so because he wants to do the right thing and he believes the right thing to do is to refrain from killing Jews.³³

Ron does the right thing but not for the right reasons: it is right not to kill Tamara, but not because she is a Jew, instead it is right not to kill her because she is a *person*. Notice that Ron is well-motivated: he sincerely wants to do what is morally right. In that respect, Ron is very different from Kant's shopkeeper who does the right thing from selfishness. But in another regard his fault is similar, he is doing the right action, but his action is not morally worthy.

Arpaly suggests two explanations of why Ron's action is not morally worthy. First, it is merely accidental that he did the right thing in this case: "Just as in the prudent grocer case it is fortunate that the prudent action also happens to be the moral one, in Ron's case it is fortunate that favouring Jews in a certain way is moral."³⁴ Secondly, Ron's reasons for action have nothing to do with the right-making features of his action.³⁵

As we might expect, these two explanations are not equivalent. Consider a case in which it is not accidental that you act rightly, but you do not act for the right reasons. Suppose that Ron wanted to kill Tamara, asked his rabbi what to do and his rabbi, whom Ron knows to be a morally good person and trustworthy, told him that it was morally right not to kill her. Ron now knows that it is morally right not to kill Tamara. Suppose Ron then wonders why it is right, and not being very good at moral reasoning, decides that it is because she is a fellow Jew. Ron's mistaken explanation for why not killing Tamara is right does not, I think, undermine his knowledge that not killing her is right.³⁶ So we can restate the example.

The Knowledgeable Extremist

Ron is an extremist, believing that killing a person is not generally immoral but that killing a fellow Jew is a grave sin. Ron would like to kill Tamara, but he refrains from doing so because he wants to do the right thing and he *knows* (on the basis of his rabbi's testimony) that the right thing to do is to refrain from killing her.

Ron does the right thing, on the basis of his desire to do what is right and his knowledge that his action is morally right. It is not accidental, then, that he does the right thing. But he does not act well: he does not do the right thing for the relevant moral reasons, namely, that Tamara is a person.

Suppose, however, that Ron's rabbi tells him why it is right not to kill Tamara. Ron believes him and so knows that it is right not to kill Tamara and he knows why, though he does not understand why. Nevertheless, Ron does the right thing, and if asked, will say that he did it because it was right and that it was right because Tamara is a person. Surely he must now count as doing the right thing for the right reasons?

There are two ways in which Tamara's being a person could be Ron's reason for action. First, he might respond to it directly, that is, he might choose not to kill her because she is a person, without explicitly forming any moral beliefs at all.³⁷ But obviously Ron did not do that: left to himself without the intervention of his rabbi he would have killed Tamara, nor would he have responded to her being a person without being told that that fact was morally significant.

Secondly, Tamara's being a person might be Ron's reason for action in the sense that he chose the action because he thought it was right, and he thought that it was right because Tamara is a person. But Ron formed the belief that it was right not to kill Tamara on the basis of the rabbi's testimony, not on the basis of his belief that she was a person. Even when he was told by the rabbi that this was the explanation of why it was right not to kill her, he was unable to draw the conclusion himself from the reasons why it was right. He *cannot* have formed the belief that it was right not to kill in response to the reasons why it was right – he is responding to testimony, not to moral reasons, just as Claire the incompetent judge was doing.³⁸

Ron might concede that he acted because of the rabbi's testimony (and his belief that the rabbi is trustworthy and reliable), not to Tamara's being a person, but nevertheless insist that he chose his action because it was right. Is this not sufficient for his action to be morally worthy? Ron has good motivations, he wants to do what is morally right and chooses in accordance with those desires, and he has moral knowledge too, but more is required for morally worthy action: you need to act for the reasons that *make* your action right. Morally worthy actions are sensitive to moral considerations and for that you need to be oriented properly, not just in your outward actions but in your motivations, your choices and your beliefs too.

Of course, the contribution that your beliefs make to morally worthy action is not the same as the contribution of your motivations: failing to be properly oriented with regard to beliefs is very different from failing to be oriented with regard to motivations. It might not be appropriate to blame Ron for his failing, but it would be appropriate to regard him, like Claire, as not fully admirable.

Does orientation matter? It might be objected that we ought not to mind whether other people have moral understanding. What matters is how they treat us, in terms of their outward actions towards us.

From the first person perspective, from which I morally evaluate my own actions and myself, it is not at all plausible that outward action is all that matters to us or that it is all that should matter. How one thinks about things, whether one has grasped what is morally right and wrong, how one makes decisions about what to do; these are important too. And these are combined in the idea of acting well. It is perfectly reasonable to be morally concerned with moral understanding, as well as doing the right thing.

This could make it appear that it is rather self-regarding or morally self-indulgent to concentrate on morally worthy action: morally worthy action is essentially about yourself rather than how you treat other people.³⁹

There are circumstances in which outward action is particularly salient. If I am in desperate need, the reasons why someone chose to give me help simply may not matter at all to me. Suppose that through a series of terrible misfortunes, I end up homeless, ill and starving. You buy me lunch and find me a doctor and a place to stay. Obviously I am very grateful to you, and I have no inclination to inquire too closely into your reasons: do you genuinely intend to help me because of my needs, or because someone told you to do so? I am just glad of some food and a place to sleep; nothing else matters to me. But this does not mean that it is not in fact important whether or not you understand why you have reason to help, even if my circumstances are so grave that I no longer care. This is a reflection of my own dire straits: it is quite reasonable, for example, for other people to be interested in whether you acted on the basis of moral understanding. It is important to us that other people grasp, for example, that when we are in need, their reason to help is not merely that we may be able to reciprocate for them sometime, or because someone told them to do so, even though people acting for those reasons might be reliable assistants. We quite reasonably value being treated well, not simply being treated rightly.

We do tend to evaluate differently those who do the same outward action. But, it may be objected, perhaps this does not show that morality is primarily concerned with acting well, rather than doing the right thing. Considerations about understanding may be relevant to an evaluation of the agent's *character* but not her *actions*.

I have argued that moral understanding (and more broadly, orientation) is essential to good character, and that there are close links between good character and morally worthy action. I suggest that we can conceive of actions in two ways, as "outward" actions, or as "actions-for-reasons". With regard to outward actions, the honest and the selfish shopkeeper do the same thing when they both give the right change, as do Claire and Judith when they pass the same sentence, and Ron and the rabbi when they refrain from killing Tamara. But with regard to action more broadly conceived, "acting-for-a-reason", their actions are not the same. There is a perfectly good sense in which action that is based on moral reasons is different from action that is not. Each of Claire and Judith and the honest and selfish

shopkeepers can and should be evaluated negatively or positively *for what she did*, not just for the kind of person that she is.

The moral evaluation of an action understood in this broad sense cannot be separated from an assessment of the agent's reasons for action. You can respond to morality through your choices and your moral beliefs as well as your motivations and your outward actions. You can be – you should be – properly oriented with regard to your whole self and you can express this orientation in morally worthy action. Moral understanding is important not just because it is a means to acting rightly or reliably acting rightly, though it is. Nor is it important only because it is relevant to the evaluations of an agent's character. It is essential to acting well.⁴⁰

4. The role of moral testimony

I have argued that it is essential to virtue and to morally worthy action that you base your moral beliefs on your moral understanding. In many of these examples there was a tension between forming moral beliefs on the basis of moral understanding and forming them on the basis of testimony. So what is the role of moral testimony?

I will begin with the place of testimony in acquiring moral understanding. To understand why p, you need 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'

If someone tells you that p, that q and that q is why p, under the right circumstances you will know why p. It is possible that you will at the same time acquire the abilities (i)-(vi). Having been told that lies that make people are happy are not always right, for example, you will immediately see why this is so, and how to apply your grasp of these considerations to other similar cases (where for example you do not lie but fail to tell the whole truth, or where a lie is the only way to avoid a really terrible outcome). But as with many abilities you will not always or even typically acquire these abilities through testimony. Instead, you have to *practise*. Most people cannot learn to ride a bike, to perform brain surgery or to type at 100 words a minute by listening to someone describe how to do so, or reading even an exceptionally good textbook. Still less can they learn to do so by having someone merely

describe the outcome at which they are aiming. These are difficult skills to master, and most people need to practise them.

Moral understanding involves intellectual rather than practical abilities, but the skills in question are difficult to acquire. Moral questions are often hard. Which actions are morally right and wrong is not always obvious, nor is it obvious why they are right or wrong, nor is it straightforward to make moral judgements in unfamiliar circumstances.

Most people cannot learn to make correct moral judgements by listening to those who are good at making judgements describe what they do, and they certainly cannot learn to do so simply by hearing what the correct moral judgements are. To get started making moral judgements yourself, of course, you need a stock of true moral beliefs on which to draw while making further judgements, and it is very plausible that to begin with, you will need to form some moral beliefs by trusting what others tell you. Recall Sarah:

Sarah is five years old. Her parents have told her not to lie, that lying is naughty and wrong. She accepts what they say.

It is obviously right for Sarah to accept what her parents say. Most young children are immature moral agents, which is why it is perfectly in order for them to trust moral testimony and defer to the judgements of the adults around them. In fact, it is likely to be impossible – or at least exceptionally difficult - for them to gain moral understanding unless they begin by forming beliefs in this way. The same is true for any immature moral agent – anyone who does not have the requisite stock of true moral beliefs - so it may be acceptable for some adults to defer to moral testimony too, if they come to judge that their starting points are too far off the mark.⁴¹ But most people cannot acquire the ability to make moral judgements themselves whilst deferring all the time to others. They need to practise drawing moral conclusions (this action is morally right, I have moral reason not to perform that action ...) and offering explanations for those conclusions (the action is right because it is helping someone in need; I have moral reason not to do that because it will make people unhappy...). Only by trying to make their own mind up about difficult moral questions will they develop the ability to make their own moral judgements, without wholly relying on the judgements of others.

This marks a key difference between knowledge (including knowledge why p) 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 to testimony as knowing why p. In section 2 I drew an analogy between moral understanding and know how and it is also true that knowing how is often not successfully transmitted by

testimony – most people need to practise riding a bike, they cannot do so on the basis of a verbal description – in Stanley and Williamson’s terms, they may know a way to ride a bike, but they do not know it “under the practical mode of presentation”. This practical mode is often not transmissible by testimony. Usually we do not even try to give someone else the abilities to put the explanation of p into her own words, or to draw the conclusion that p from q, for example, by testimony. Further testimony of any kind will not usually help you to understand why p. Therefore *even if moral understanding is a species of propositional knowledge after all*, it will not have the same relationship to testimony as other forms of propositional knowledge.

Suppose that you are a mature moral agent and you have no reason to think that moral understanding is out of your reach. Should you *ever* trust testimony about moral matters? Testimony can still be important in a number of ways.

In the first place, understanding why an action is morally right or wrong often involves knowing a combination of both moral and non-moral matters. Understanding why you should give money to X, for example, involves both truly believing that X is in need and understanding why you should help those in need. It is perfectly acceptable to trust the testimony of someone on the non-moral matter – that X is in need – and use this in forming your moral judgement. So testimony can play a very significant role. What is more, often it is difficult to disentangle the non-moral and the moral matters and in practise there may be no clear division between the two. Consider the following example, from Julia Driver:

It might make perfect sense, for example, for someone to delegate a decision to an expert. Suppose that Stanley must decide whether or not to authorize keeping his ailing father alive using “extraordinary” means. Suppose that Stanley has good reason to trust another person – someone with experience, perhaps the family’s physician. Then, when the physician says “Trust me, Stanley, you ought not to do it – it will cause unnecessary pain”, doesn’t it make sense to say that Stanley knows it (given it is true and he believes it)?⁴²

What exactly does Stanley learn from the family doctor? The physician is clearly a medical expert, and can speak confidently about the probable outcomes of various surgical procedures. He may be able to tell Stanley that his father would be suffering unnecessarily, in the sense that without medical intervention, he would die peacefully and without pain. Stanley obviously has very strong reason to trust what he says, given the doctor’s expertise. But the doctor may alternatively be saying to Stanley that the suffering would be unnecessary in the sense that his father’s life is no longer worth living, and it ought not to be prolonged any further. Since the doctor is so experienced, Stanley would be wise to listen to his advice about

the value of his father's life, think the question through himself and see if he agrees with the doctor. But it is much less obvious that Stanley has a strong reason to simply place his trust in what the doctor says on this moral question, and defer to his judgement. At the same time, it is clearly difficult to separate completely the non-moral and moral issues. It may be unclear to both what kind of judgement Stanley is asking for from the doctor.

I said at the beginning that taking moral advice was perfectly acceptable even though trusting testimony is not, and we are now in a position to see why, by distinguishing between two responses to testimony. First, you may *trust* or *defer* to moral testimony, where you simply believe what is said to you. You make no attempt to gather the reasons why p and draw conclusions yourself, or to devise explanations for moral propositions that you have accepted. You simply believe what you are told.

It does not follow that you have suspended use of your critical faculties entirely, believing any moral proposition that is said to you by anyone. You may take into account how expert and trustworthy the speaker is, rating her as a source of moral testimony: one who is usually right, or who has the relevant experience to be right about this kind of issue most of the time. Nevertheless, once you have decided that she is reliable (or at least not wholly unreliable) you simply trust what she says without exercising your own judgement on that particular matter. In this case, you are relying on the judgement of other people.

Recall Eleanor, who was perfectly capable of working out for herself whether it was wrong to eat meat, but chose to trust her friend's testimony instead. If Eleanor fails to use her moral understanding, she is not properly oriented and she cannot act well. To act well on the basis of moral beliefs, it is not sufficient that you have moral understanding, you need to *use it* to form your moral beliefs too, so that your belief that your action is right, for example, is explained by the reasons why it is right. Trusting testimony is a *rival* basis for your belief. At best it is worthless, since you need to base your belief on another ground anyway, and at worst it is harmful, since it could make it less likely that your belief is grounded in the right way, hence make it less likely that you act well.

Secondly, you may treat the testimony as *moral advice*, which you subject to critical scrutiny and you decide whether or not to accept *on its own merits*. You take into account what others have said to you as a guide to your own reflections. Moral advice is extremely important, as moral questions are often difficult to decide, not just because they are inherently hard but also because your own desires, interests and emotions can bias you and lead you astray. Advice from others, who can put forward another point of view, make salient the interests of others and try to help you to see more clearly, is often essential to your gaining genuine moral understanding. But of course, in treating moral testimony as advice, even though you are paying attention to the testimony, you are not simply putting your trust in it or deferring to it, you are using your own judgement about the matter at issue. And it is not at all

unusual for those offering advice to make explicit that they expect you to respond in this way rather than with deference, by prefacing their advice with remarks such as, for example, “look at it this way...” or “have you considered this point of view...?”

Of course, in practise it may be difficult to distinguish these two ways of treating moral testimony, but nevertheless in principle the difference is clear. If you are trying to develop and use your moral understanding, it is perfectly in order to take testimony on explicitly moral matters as moral advice to guide your own reflections. Stanley, for example, would be wise to take the advice of his doctor about the value of his father’s life. But what you cannot do is to treat testimonial evidence as having authority independently of the reasons why *p* that you acknowledge or to which you can gain access. There are significant limits to the possibility of a “division of epistemic labour” if you want to acquire and use moral understanding: there are no benefits to the theft of other people’s conclusions over honest toil on your own.

So far I have assumed that you can base your moral beliefs on your moral understanding and I have argued that you should. But suppose that you cannot. If your moral judgement about a particular matter is bad and does not improve when you reflect, you may never have moral understanding. You also have a serious problem, however excellent your judgement, if you do not have any idea of the reasons why *p*. If you are in this position, it is simply not possible for you to gain moral understanding.

If moral understanding is unavailable for you, there is no point in your trying to acquire and use it. Given your situation, morally worthy action and proper orientation with regard to your moral beliefs are out of reach. But if you are lucky and you have access to reliable people whom you can trust, you might be able get moral knowledge and as a result, do the right thing. Since doing the right thing is very important, *you should trust moral testimony from trustworthy and reliable sources and defer to moral experts if you cannot gain moral understanding*. Claire the incompetent judge and Ron the extremist may be in this position. They cannot acquire and use moral understanding, and they cannot act well. So it is better that they trust the testimony of some reliable person – Judith and the rabbi respectively – thereby gaining moral knowledge and doing the right action.⁴³

What if you are simply not sure whether you can acquire moral understanding, but there are other people whose moral testimony you could use? There probably are situations of this kind in which it is reasonable for you to trust moral testimony, in order to raise the chances that you do the right thing. But whether it is right to trust testimony in these circumstances depends on a number of factors that will be different in different situations, in particular: (1) how likely you are to acquire moral understanding if you try; (2) how likely you are to gain moral knowledge if you trust the testimony of the people around you; (3) the

relative importance of morally worthy action and, more generally, proper orientation, compared to doing the right thing but not for the right reasons.

It is difficult to assess the probability that you will acquire moral understanding, if you tried, and the probability that you would gain moral knowledge, if you trusted the people around you. As we saw right at the start, it is a matter of controversy who are the moral experts (and even if there are any) and there is a lot of disagreement about many moral questions. So it is far from obvious whom to trust about moral matters. There will be situations in which you are more likely to gain moral understanding by thinking things through yourself than you are to gain moral knowledge by trusting others. Moreover, I have argued that it is reasonable to regard morally worthy action as more valuable than simply right action (that is not for the right reasons). It is not always worthwhile trying to develop moral understanding; sometimes it may be better to trust the judgement of others.⁴⁴ But on the other hand, in many circumstances it is reasonable for you not to trust moral testimony, but to take advice and use your own judgement, in order to leave open the possibility at least of your acquiring and using your moral understanding, orienting yourself properly and acting well.

5. Conclusion

According to optimists, you should trust moral testimony (whenever the speaker is knowledgeable and trustworthy), and if you do so you can acquire moral knowledge. According to pessimists, normally you have good reasons not to form your moral beliefs on the say-so of others. I have accepted that optimists are right that you can gain moral knowledge by deferring to moral experts and putting your trust in moral testimony. Where optimists go wrong is in their assumption that all that we are or should be interested in is moral knowledge. I have argued that we should be trying to gain moral understanding, in which case pessimists are right – at least in circumstances where moral understanding is available - that you have good reasons not simply to trust moral testimony.

If you want to be properly oriented to moral reasons, you have excellent *instrumental* reasons not to defer to experts or to put your trust in moral testimony. But these are practical reasons not epistemic ones, so it does not follow that it is epistemically rational to treat moral testimony in this way. Of course, these are unusual practical reasons, because they are based on a goal – moral understanding – which is (arguably) *epistemically* valuable. But it is a separate question whether not placing your trust in moral testimony or deferring to moral experts is epistemically rational. I think that it can be epistemically rational, but unfortunately this is not an issue that I can pursue here.⁴⁵

Instead I want to finish with a few brief remarks about moral philosophy. Once we recognize that moral understanding is important in ethics, we also see that the role of moral

experts is limited. There is no point merely deferring to an expert when you are trying for understanding rather than knowledge. But this is not to say that there is no room for experts at all. Some people may be better informed than others about moral matters, have more experience and better judgement. Experts can still play an important role in helping others achieve moral understanding, by explaining as far as possible what kinds of considerations are important, how important they are and in what way. They can give advice. Moral philosophers could have a role as moral experts, if a limited one.

Can moral philosophy help us acquire moral understanding? Arguably this is one of the aims of moral philosophy, and as such it is reflected in our practice. We do not typically defer to others, even those who have studied certain moral problems for longer than ourselves. Through moral philosophy, we can become very good at devising moral theories, arguments for and against those theories, and so on. But one can be very skilled at arguments and constructing moral theories without having or developing good moral judgement. Do moral philosophers tend to have judgement that is as good as or better than anyone else?

Moral philosophy is characterized by sustained reflection on questions of the content and authority of morality. It involves taking rival views seriously, thinking about unusual examples and thought experiments. This sort of reflection could improve your moral judgement. It might bring to your attention considerations that you had ignored or undervalued. You could correct casual or biased first impressions. Could this kind of reflection be corrupting, making your judgement worse not better? It is possible for instant judgements to be better than reflective ones. Reflection can introduce bias, mistakes of reasoning, the influence of others can lead you in the wrong direction, and so on. Reasoning about ethics can turn into a rationalization of whatever it is most convenient for you to think.

These questions about moral philosophy are troubling and should concern us. But on the other hand they do not give us reason to think that reflection of the kind characteristic of moral philosophy - aiming to counteract bias and mistakes and so to help us appreciate moral reasons - typically introduces more of these than it corrects. So in the absence of further reasons to doubt the outcomes of philosophical reflection, we are not compelled to conclude that moral philosophy leads us astray. And we have some grounds for the hope that through moral philosophy we may to some (perhaps limited) extent, gain genuine moral understanding.

* I would like to thank John Broome, Jessica Brown, Roger Crisp, John Hawthorne, Elijah Millgram, Duncan Pritchard, Tim Williamson, audiences at the universities of Sheffield, Bristol and Oxford and the anonymous referees from this journal for helpful comments on earlier drafts of this paper.

¹ By “moral testimony” I mean testimony with *explicitly moral content*. This might be testimony about what is morally right or wrong: “It is morally right to help those in need”. Or it might be testimony using “thick” ethical terms, such as: “It is courageous to face danger without flinching.” Or it might be about our moral reasons for action: “I have moral reason not to lie.” I will not try to give a definition of what it is for the content of a proposition to be explicitly moral, but I will assume that we can, at least roughly, distinguish between these propositions and others that have what I will call *ordinary non-moral factual content*. Of course there is also a controversy about ordinary testimony, namely whether testimony is a basic source of justification (“non-reductionism”) or not (“reductionism”), discussed for example in Jennifer Lackey and Ernie Sosa, eds. *The Epistemology of Testimony* (Oxford: Oxford University Press, 2006). But this debate can be put to one side. If Eleanor was learning ordinary non-moral facts from her friend, both reductionists and non-reductionists would regard her as justified (for she knows that her friend is trustworthy and reliable).

² Those who think we can gain knowledge from moral testimony include Karen Jones, “Second-hand moral knowledge” *Journal of Philosophy* 96 (1999): 55-78; Julia Driver, “Autonomy and the Asymmetry Problem for Moral Expertise,” *Philosophical Studies* 128 (2006): 619-644; and Robert Hopkins, “What is wrong with moral testimony”, *Philosophy and Phenomenological Research*, 74 (2007): 611-634. Elizabeth Fricker, in her “Testimony and Epistemic Autonomy” in J. Lackey and E. Sosa *ibid.* is also optimistic. Those who disagree include R. P. Wolff “The Conflict Between Authority and Autonomy” in *Authority* ed. Joseph Raz (Oxford: Blackwell, 1990). The terms “optimist” and “pessimist” are from Hopkins, *ibid.*

³ Bernard Williams, *Making Sense of Humanity* (Cambridge: Cambridge University Press, 1995) p. 205.

⁴ However, as Hopkins points out, some moral truths are trivial, and some non-moral matters (will the rope hold my weight? Will the bridge collapse if I stand on it?) are extremely important. So this cannot explain why in general deferring to experts and trusting testimony is more acceptable for ordinary factual matters than for moral questions (see Hopkins, “What is wrong with Moral Testimony” pp. 621-3).

⁵ Hopkins makes this argument for gaining moral knowledge by testimony (Hopkins *ibid.*, pp. 623-6).

⁶ Note that the assumption that there is moral knowledge is compatible with moral realism, the theory that there are truths in ethics that are not dependent on our beliefs and attitudes about them. But it is also compatible, I believe, with a wide range of other antirealist theories, including sophisticated forms of subjectivism and non-cognitivism, quasi-realism and moral fictionalism (though they may have to give a revisionary account of what moral knowledge is).

⁷ The defence is given by Jones “Second-hand moral knowledge”, Driver “Autonomy and the Asymmetry Problem for Moral Expertise,” and Hopkins “What is wrong with Moral Testimony?”

⁸ Hopkins *ibid.* addresses the puzzle about moral testimony and suggests that the best explanation is that though moral testimony can give us knowledge, there are reasons (perhaps moral reasons) why we ought not to use it. In some ways, this paper is an elaboration of that suggestion.

⁹ It is important to note that this conception of moral understanding – understanding why p (where p is some explicitly moral proposition) – differs from some other types of understanding. It is not the same as understanding the proposition itself, for example, understanding what “it is morally right to help those in need” means. I assume that you need to understand the proposition itself to believe the proposition, let alone to know it or to understand why it is the case. My conception of understanding why p (where p is some moral proposition) should also not be confused with the idea of understanding morality as a whole. I do not have anything to say here about what it is to understand morality (considered as a whole subject area), though it may include having some moral knowledge or some moral understanding in my sense. My conception of moral understanding should also be distinguished from understanding that p (even where p is a moral proposition), a phrase that I think is often used as equivalent to believing that p (or perhaps knowing that p). It is also important to note that I am giving an account here specifically of moral understanding, not of understanding generally.

¹⁰ Other philosophers who have written about understanding in general (rather than moral understanding specifically) have taken a different view. For example, Linda Zagzebski, “Recovering Understanding” in *Knowledge, Truth and Duty: Essays on Epistemic Justification, Responsibility and Duty* ed. M. Steup (Oxford: Oxford University Press, 2001) develops a conception of understanding according to which understanding is transparent – if you understand that p, you understand that you understand that p - and which is not factive: you can understand that p even though p is false. This is completely different from the concept of moral understanding in which I am interested, and it is also, I think, rather different from our everyday use of the term. Jon Kvanvig, *The Value of Knowledge and the Pursuit of Understanding* (Cambridge: Cambridge University Press, 2003) has a concept of understanding more similar to mine (though my account of moral understanding is developed in ways with which he might not agree, especially with regard to the importance of certain kinds of cognitive ability). I would like to emphasize that moral understanding as described here is compatible with a variety of metaethical theories, including moral realism, but also with various forms of anti-realism (including sophisticated subjectivism, quasi-realism and fictionalism), provided that they include a conception of truth in ethics.

¹¹For example, Philip Kitcher, “Scientific Knowledge”, *Oxford Handbook of Epistemology*, ed. P. Moser (Oxford: Oxford University Press, 2002); James Woodward, *Making Things Happen: A Theory of Causal Explanation* (Oxford: Oxford University Press 2003) p. 179; Peter Lipton, *Inference to the Best Explanation*. (2nd edition) London: Routledge, 2004) p. 30. Stephen Grimm (“Is Understanding a Species of Knowledge” *British Journal for the Philosophy of Science*, 57 (2006): 515-535) defends a view of understanding according to which it is a “species of knowledge”, that involves a “psychologically thick” notion of belief: if you understand that p, you must have “grasped” that p. He

does not discuss issues such as accepting that p on the basis of testimony, so it is not clear whether or not he would agree that understanding differs from knowledge in that regard.

¹² This testimony might be the trigger that Eleanor needs to think through the matter herself. She might as a result grasp why the welfare of animals is important, in such a way that she could make similar judgements about similar cases. If so, then she would understand why p.

¹³ A very similar case is presented by Duncan Pritchard, “Knowledge, Understanding and Epistemic Value” in *Epistemology*, ed. Anthony O’ Hear, Cambridge: Cambridge University Press, forthcoming), and he draws the same conclusion, that understanding differs from knowledge that and knowledge why.

¹⁴ Note that moral understanding does not require a grasp of reasons “all the way down”: this would be far too demanding. Nor does it require a grasp of philosophically sophisticated justifications or of moral theory. It does not need any more than the kind of grasp of reasons that most of us have when we realize that we should help others, take care of our friends and family, and so on.

¹⁵ Consider the different grounds that you might have for believing a moral proposition p, for example “Action X is wrong”. You might base your belief that action X is wrong on the grounds that X is a lie and that lying to others is wrong because it fails to respect them, for example. These factors are reasons why it is wrong, that you could use to *explain why* that action is wrong. By contrast, you may have evidence that is *not* of the reasons that make the proposition true, that could *not* be used to explain why it is true. One obvious source of such evidence is testimony. If a reliable person tells you that p, or even if you discover that they believe that p, this is evidence for p; but it does not give you a clue why p is the case. You cannot offer as an explanation of why action A is wrong: my friend told me it was. We can call this type of evidence, that is not of reasons why p is true, *non-explanatory evidence*. The argument here about the limits of testimonial evidence to moral understanding generalizes, I think, to other forms of non-explanatory evidence (such as the evidence of whether or not other people believe that p, including whether one’s “epistemic peers” believe that p) and has significant implications for the appropriate response to moral disagreements, for example. I discuss these implications further in my forthcoming book, *The Beloved Self* (Oxford: Oxford University Press).

¹⁶ This is perhaps one of the reasons why it seemed plausible to some that the virtues are unified in the sense that if you have one virtue, you must have them all. It is not possible properly to understand one isolated moral reason. But at the same time, you do not have to grasp the nature and significance of all moral reasons in order to understand one of them.

¹⁷ I do not mean to suggest here that moral thought has to take place in isolation. Thinking about moral questions with other people can be a way of coming to moral understanding. But if each individual is to have moral understanding, each must grasp the connections between p and the reasons why p, even if they have acquired that grasp through collective deliberation.

¹⁸ I leave it open whether having moral understanding is identical with having these abilities, or whether they are connected in some other way that is not identity – but where having moral understanding requires that you have these abilities (to some extent at least).

¹⁹ 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). Understanding is often associated with certain sorts of feeling: a flash of enlightenment; a light dawning. But these are not necessary: you need have no particular feelings at all when you finally come to understand why X is morally wrong. Nor are they sufficient. However much you feel the light dawning as you confidently turn aside from the needy, you do not understand - you cannot understand - that it is right to do so.

²⁰ For example, Elizabeth Anscombe argued that dropping the atomic bomb on Hiroshima was wrong because intentionally killing innocent civilians in war is absolutely wrong (Elizabeth Anscombe, "Mr Truman's Degree," in *The Collected Philosophical Papers of G. E. M. Anscombe, vol. III, Ethics, Religion and Politics* (Oxford: Blackwell, 1981). Let us assume that she was mistaken about that: intentionally killing innocent civilians in wartime can be permissible provided that the benefits are sufficiently great. Let us also assume that in Hiroshima the benefits were not sufficiently great. Given these (controversial) assumptions, Anscombe was right to condemn dropping the bomb, and she partially understood why doing so was wrong: that is involved the killing of civilians was certainly relevant. But she did not fully understand why it was wrong, though she believed that she did.

²¹ This is the same sort of case as presented by Pritchard ("Knowledge, Understanding and Epistemic Value") for a similar conclusion. Pritchard concludes that understanding, but not knowledge, is compatible with "environmental" epistemic luck, though he denies that it is compatible with all forms of epistemic luck. I am inclined to go further and to say that even if your own text-book was extremely inaccurate, but that you read in it the truth that Stalin killed millions, you could as a result understand why he was an evil person, though obviously you could not know that he was evil or know why on that basis (I think that Pritchard would deny this). I am inclined to think that you have moral understanding because your belief that Stalin killed millions is true and you formed it reasonably (assuming that you had no reason to think that your textbook was generally inaccurate) and you have grasped correctly he was evil because he killed millions - and this is all that is required for understanding why he was evil.

²² The best known defender of the claim that know how is *not* propositional knowledge is Gilbert Ryle, "Knowing How and Knowing That", in *Gilbert Ryle: Collected Papers, Volume 2*. New York: Barnes and Nobles (1971): 212-225. The thesis is also endorsed by David Lewis, "What Experience Teaches", in ed. W.G. Lycan, *Mind and Cognition: A Reader* (Oxford: Blackwell, 1990): 499-519. Recently, the opposite claim - that know how is a form of propositional knowledge - has been forcefully argued by Jason Stanley and Tim Williamson, in their paper "Knowing How" *Journal of Philosophy* 98.8 (2001): 411-44. The argument given by Stanley and Williamson is complicated and I cannot do justice to it here. They claim that knowing how is a species of propositional knowledge, specifically that, for example, "Hannah knows how to ride a bicycle" is true relative to a context c "if and only if there is some contextually relevant way w such that Hannah stands in the knowledge-that relation to the Russellian proposition that w is a way for Hannah to ride a bicycle, and Hannah entertains this proposition under a practical mode of presentation." (Stanley and Williamson, *ibid.*, p. 426).

²³ Stanley and Williamson (*ibid.*) press the point that having knowledge how can involve having certain complex dispositions, and yet knowledge how may still be a species of propositional knowledge, so they would not find convincing the first reason I give for thinking that moral understanding and propositional knowledge are distinct. But even if they are right about that, it is still true that moral understanding seems to be compatible with epistemic luck in a way that propositional knowledge is not, in which case understanding why p (where p is a moral proposition) cannot be a species of propositional knowledge.

²⁴ Tim Scanlon, *What We Owe To Each Other* (Cambridge Mass: Harvard University Press, 1998), pp. 147-60. According to Scanlon, the justification of your action must be one that others could not reasonably reject.

²⁵ It may be very unlikely that anyone will ever ask you to justify what you do. You may know that no one will ask you: perhaps you acted in secret; maybe everyone is more concerned with their own affairs. Why should it be important that you can justify yourself to them, if you know you will never have to? Perhaps you find yourself imagining other people watching you and you imagine trying to justify yourself to them. You might feel very uncomfortable if you found yourself unable to explain to them in your own words why what you were doing was right. But if you preferred not to exercise your imagination in this way, it is difficult to see why you should care about justification to others at all, in situations where no one wants a justification.

²⁶ Compare Aristotle: “the “product” is brought to completion by virtue of a person’s having wisdom and excellence of character, for excellence makes the goal correct, while wisdom makes what leads to it correct” (Aristotle, *Nicomachean Ethics*, trans. Sarah Broadie and Christopher Rowe (Oxford: Oxford University Press, 2002) 1144a6-8). I want, as far as possible, to leave open the connection between the two – for example, I do not intend to take a stand on whether, if you have good judgement, you need a separate motivation in order to act – or whether you must have a good motivation, if your judgement is good.

²⁷ John McDowell defends a similar view, though he says that the virtuous have a “reliable sensitivity to a certain sort of requirement which situations impose on behaviour” (John McDowell, “Virtue and Reason” in *Virtue Ethics* eds. Roger Crisp and Michael Slote (Oxford: Oxford University Press, 1997, first published in *The Monist*, 62, (1979): 331-50); p. 142). He regards this sensitivity as delivering knowledge, though as I have explained, I think it is better characterized as moral understanding. There are other important features of McDowell’s conception of virtue on which I do not want to make a stand (for example, his view that this “knowledge” is sufficient for action).

²⁸ In order to have some virtues, you may not need to form explicitly moral beliefs. It may be sufficient to be generous, for example, that you simply judge that someone else is in need and that you could help them, and you go on to do so. But it is likely that you will have to make explicitly moral judgements some of the time, and with regard to some virtues – like justice, for example – you will almost certainly need to judge outcomes in moral terms (in this case, as just or unjust). There is in fact no sharp distinction between moral and non-moral motivations, and moral and non-moral judgements. For example, one of the important ways that we recognize moral reasons is through emotional responses,

some of which involve a distinctively moral attitude, or a combination of non-moral response and moral judgement: we may feel anger at injustice or at cruelty; admiration at an example of courage and anxiety at whether we can live up to the same standards; respect for others can be a feeling which is not explicitly moral (which does not involve explicitly moral beliefs) or a distinctively moral attitude towards others that we deliberately take up. There is often continuity between our non-moral attitudes and emotions and those that involve explicitly moral descriptions and lead us to form explicitly moral beliefs. An honest person, for example, grasps why honest actions are important and why she ought to perform them. A courageous person is motivated to face danger, precisely when she grasps that it is worth doing so. Much (though not all) of the motivations and judgements associated with the virtues are explicitly moral.

²⁹ For example, Claire's belief that A is right is not counterfactually dependent on R or her belief that R, but on what Judith says. Of course, it will be difficult to distinguish these different dependences in practise, since in fact what Judith says about whether A is right will itself depend on R. But though Judith is virtuous and trustworthy, she is not infallible, and Claire's belief depends on Judith's testimony, not on R (or her belief that R), when the two come apart. Counterfactual dependence is only a sign of the basis of one's belief though and it is not an infallible sign. If Claire gains and uses moral understanding, for instance, her belief that A is right may be based on *both* Judith's testimony *and* on her belief that R; both of these explain why she has that belief. I take it that it is sufficient for proper orientation is that one explanation for your moral belief is your moral understanding, it may have other explanations too. When Claire does not have or does not use moral understanding, by contrast, her belief that A is right is explained by Judith's testimony, not by her belief that R.

³⁰ As Aristotle says, the virtuous person "sees the truth in each case, being as it were a standard and measure of them" (*Nicomachean Ethics* 1113a32-3).

³¹ Immanuel Kant, *Groundwork of a Metaphysic of Morals* trans. H. J. Paton (London: Routledge, 1991; page references cite the volume and page number of *Kants gesammelte Schriften* (published by the *Preussische Akademie der Wissenschaften*, Berlin: W. de Gruyter, 1902-), 4:397. What is it to act for the right reasons? Kant himself thought that morally worthy action was doing your duty *from duty* – though it is not exactly clear what that means. One possibility is that to do the morally right act for the right reasons, you must choose it under an explicitly moral description "the morally right action" or "the action which I have overriding moral reason to perform" or "my duty". But that is a mistake. The use of the concept of a duty or of moral rightness - or indeed the use of any explicitly moral concept - is not essential to morally worthy action. What matters is that you respond to moral reasons, not that you do so under an explicitly moral description. In fact, as Williams emphasizes, we sometimes expect people to respond to moral reasons without thinking of them in explicitly moral terms, and responding to them under an explicitly moral description "it's my wife; and in situations like this it is morally acceptable to save one's wife" seems to be "one thought too many" (Bernard Williams, *Moral Luck* (Cambridge, Cambridge University Press, 1981), p. 18). Of course Williams's argument is compatible with moral beliefs playing an important background role, for example, your recognition that it is morally acceptable to save your wife might influence what you do even though it does not consciously

cross your mind. But it is likely that at times when you act, you will form some moral beliefs and they will play an important role. Sometimes there is no clear distinction between explicitly moral and non-moral considerations. Sometimes non-moral motivations lead you astray and you need to form moral beliefs to correct them. Sometimes what is morally required is complicated and you need to deliberate and think in explicitly moral terms in order to work out what is morally right in your circumstances. It does not follow, of course, that our forming explicitly moral beliefs does actually help us to respond to moral reasons. Our moral beliefs can fail to fulfil this function in (at least) two ways: first, they may be false, and lead us astray; secondly, we may fail to act on them. And of course, they can fail in both ways at once. But if you do act on your moral beliefs, I will argue, you need to have and to use moral understanding in order to act well.

³² Arpaly's final account of morally worthy action involves sensitivity to the right-making features of an action – plus deep moral concern – which is similar to my conception of the two aspects of orientation (Nomy Arpaly, *Unprincipled Virtue*. Oxford: Oxford University Press, 2003), p. 84.

³³ Arpaly *ibid.* p. 74. Arpaly describes Ron as believing deeply that killing a person is not generally immoral but that killing a fellow Jew is a grave sin.

³⁴ Arpaly *ibid.* p. 74.

³⁵ Arpaly *ibid.* p. 74.

³⁶ Of course if he believed that killing Tamara was wrong on the basis of a false belief, he would not have knowledge. But he does not believe it on that basis, he believes it because he was told it by a source he knows to be trustworthy, and it is hard to see that his justification for that belief is undermined because he has independently come up with a mistaken explanation for it. Compare: you have a lot of knowledge from perception, and sometimes you come up with mistaken explanations for it, e.g. you know that the sun rose in the morning because you saw it, even though you mistakenly believe that it did so because the sun orbits the Earth. The mistaken explanation does not undermine your perceptual knowledge.

³⁷ He would then be responding to moral reasons *de re* (rather than under an explicitly moral description).

³⁸ Ron's action is therefore counterfactually dependent on the rabbi's testimony (specifically, his testimony that the action is right) together with his belief that the rabbi is reliable and trustworthy, but not counterfactually dependent on his belief that Tamara is a person. Of course, this is complicated by the fact that the rabbi also testifies that Tamara is a person and that this fact is morally relevant, but if the rabbi had given a different explanation of the rightness of the action, Ron would still have done what the rabbi told him was right, whatever he believed about Tamara's being a person. This is, I think, grounds for claiming that Ron's belief that Tamara is a person is not among the causes of his action, and therefore is not among his reasons for action. It is part of what is sometimes called the "standard model" of acting for reasons, defended for example by Donald Davidson ("Actions, Reasons and Causes" in his *Essays on Actions and Events* (Oxford: Oxford University Press, 1980), pp 3-19, especially p. 9), that acting for the reason R is being motivated to act by the belief that R. The standard model has recently come under attack on the grounds that your action can be caused by a belief and a

desire without your acting for reasons (David Velleman, *The Possibility of Practical Reason* (Oxford: Oxford University Press, 2000), pp. 2-9). I am sympathetic to these concerns, but I think that a causal connection between your belief that R and your action is going to be part of any plausible account of acting for the reason R, even if it is not sufficient: in order for a reason for an action to be *your* reason for action, it must *cause* your action. Kieran Setiya (*Reasons without Rationalism* (Princeton, NJ: Princeton University Press, 2007) pp. 28-48), however, suggests that though this is true in the “overwhelmingly typical” case, it may not be true in unusual cases, such as in self-deceptive action. He suggests a different connection: “In what we may think of as the paradigm case of acting for a reason, an agent takes it that he is hereby doing ϕ because he believes that p, where this is the “because” of motivation” (Setiya, *ibid*, p. 45). In other words, to act for the reason p, you must believe that your belief that p caused you to act (even if in fact it did not). Circumstances in which the belief that p did not in fact cause you to act (though you think that it did) Setiya describes as “marginal” cases of acting for a reason. But I think a better description of this kind of case is: you thought that you acted for the reason that p but in fact you did not – since your belief that p did not in fact cause you to act. In any case, Ron has no reason to think that his belief that Tamara is a person caused him to act (either directly or indirectly, via a belief that it is right not to kill her). So even if Setiya’s account of acting for reasons is correct, Ron has no grounds for claiming that he is acting for the right reasons (and if my interpretation of what Setiya calls the “marginal” case is correct, he does not act for the right reasons).

³⁹ Arpaly argues that the moral worth of an action corresponds to the moral praiseworthiness or blameworthiness of the agent, that is, the extent to which she deserves praise or blame for performing those actions (Arpaly, *Unprincipled Virtue*, p. 69). If this is right, there is another way in which focussing on morally worthy action could be self-indulgent: the agent could primarily be motivated to be praised. It would not, of course, be particularly admirable to aim to perform morally worthy actions because you want praise (though of course Arpaly does not suggest that it is). But according to my conception of acting well, morally worthy actions are right actions performed for the right reasons. Since *to get praise* is not usually the right reason for action, someone who was motivated primarily by praise would not be performing morally worthy actions. Instead, they would choose the right because there were forceful moral reasons to do so. Acting for these moral reasons is not self-serving or self-indulgent.

⁴⁰ This is, I think, similar to Aristotle’s view that the end of *phronesis* is not the production of something distinct from the productive process but “doing well serves itself as end” (*Nicomachean Ethics* 1140b7-8).

⁴¹ Children may acquire moral understanding by learning the habit of doing what is right and not what is wrong, and eventually coming to grasp the reasons why some actions are morally right and others wrong (this is a simplified version of Aristotle’s conception of moral education). At the first stage of this acquisition, they learn from the testimony and example of those around them what they have moral reason to do (and not to do) – so moral testimony plays a crucial initial role – then they habitually do the right thing and testimony plays less of a role. The final stage of acquiring moral understanding, however, where they come to grasp why it is right to do what they already habitually do, requires them

to exercise their own judgement. It is this stage that needs practise, and moral testimony can play at most a guiding role.

⁴² Driver, “Autonomy and the Asymmetry Problem for Moral Expertise”, p. 639.

⁴³ Someone might reasonably think that moral understanding is valuable, but that since it is unavailable to many people (perhaps even to most of us), we should trust the testimony and defer to the judgements of the experts. Plato seems to express this sort of view in the *Republic*. The Guardians who have undergone the right education and training understand why justice and the other virtues are important, and are just and virtuous themselves. Ordinary people are not capable of going through the education required of the Guardians and so can never have moral understanding. It is best, Plato thinks, if society is set up with the Guardians in charge and with everyone else doing what they are told. The moral understanding that the Guardians have is clearly extremely valuable, on Plato’s view, but it is simply not open to all. Deferring to the Guardians is appropriate, not because moral understanding is not worth having, but because, however long and hard they tried, most people would never attain it. The next best option is for them to do what they are told by people who do. (It is, of course, very controversial whether Plato thought that the Guardians had moral knowledge or what I would call moral understanding. For an argument that they acquire knowledge see Gail Fine, “Knowledge and Belief in *Republic* V-VII” in her *Plato on Knowledge and Forms* (Oxford: Oxford University Press, 2003). For a reading more similar to mine, see Julia Annas, *An Introduction to Plato’s “Republic”* (Oxford: Oxford University Press, 1981)).

⁴⁴ This is the best interpretation, I think, of a well-known example from Karen Jones, of a situation in which Peter and his housemates are trying to choose someone else to share their house:

White women and women of color had advocated rejecting three white men on the basis of their perceived sexism, and in one case, racism. Peter had a settled and serious commitment to the elimination of racism and sexism, but he was not very good at picking out instances of sexism and racism. Thus, his commitment often remained theoretical and he saw fewer occasions as calling for antisexist and antiracist responses than he might have seen. Such blindness can sometimes indicate insincerity but in Peter’s case it did not. He genuinely did want to understand the perspective of the women and he wanted to be able to share it, *if* indeed it turned out to be correct. He could pick out egregious instances of sexism and racism, and could sometimes see that “sexist” or “racist” applied to more subtle instances when the reasons for their application was explained to him, but he seemed bad at working out how to go on to apply the word to nonegregious new cases. Problems arose for Peter when he could not see the reasons why the women were calling someone sexist, and could not see, or could not see as evidence, the considerations that the women thought supported viewing the would-be members as sexist... I think that in this case, Peter should have been willing to accept the women’s testimony that these men were sexist. But for him to have done so would have been to accept and act on a moral judgement on someone else’s say-so. It would have been to borrow moral knowledge. (Jones, “Second-hand moral knowledge”, pp. 59-60).

Peter has evidence that he has poor judgement in this area and that despite his best efforts he has been unable to acquire moral understanding. He also has evidence that in trusting the women's testimony he would acquire moral knowledge. In this situation, it may be right for him to do so.

⁴⁵ I defend the claim that epistemic rationality for moral beliefs is defined in terms of moral understanding in my *The Beloved Self*.