


ARCHIVES

Praise without Perfection

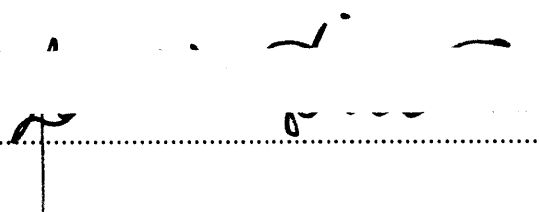
by
Paulina Anna Sliwa
MPhysPhil
Oxford University, 2007



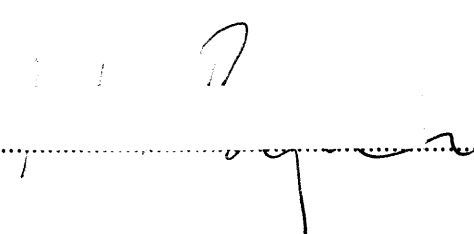
Submitted to the Department of Linguistics and Philosophy in Partial Fulfillment of the Requirements for the Degree of Doctor of Philosophy in Philosophy at the Massachusetts Institute of Technology.

September 2012
(c) 2012 Paulina Sliwa. All rights reserved.

The author hereby grants to MIT permission to reproduce and to distribute publicly paper and electronic copies of this thesis document in whole or in part in any medium now known or hereafter created.

Signature of Author: .....
Paulina Sliwa
July 19th 2012

Certified by:.....
Richard Holton
Professor of Philosophy
Thesis Supervisor

Accepted by: .....
Alex Byrne
Professor of Philosophy
Chair of the Committee on Graduate Students

Praise without Perfection

by

Paulina Anna Sliwa

Submitted to the Department of Linguistics and Philosophy on July 19th 2012 in Partial Fulfillment of the Requirements for the Degree of Doctor of Philosophy in Philosophy.

Thesis Abstract

This dissertation is about agents like us, who are both epistemically flawed and morally imperfect. First, how should such agents form and revise their moral beliefs? Second, how should we morally evaluate the actions of such agents – under which conditions do their actions have moral worth? I approach these questions by focusing on moral testimony. Moral testimony is particularly interesting in this context because we often rely on it to compensate for our moral and epistemic imperfections. Yet it has been viewed with suspicion by moral philosophers and epistemologists alike. I argue that these suspicions are unfounded: moral testimony is an important source of moral knowledge – in some situations it is the only source of moral knowledge available to agents like us – and actions on the basis of moral testimony can have moral worth. In arguing for this conclusion, I develop a novel account of moral worth. I argue that moral knowledge is central to moral worth. The epistemic question about how we should form our moral beliefs and the moral question about when our actions are morally praiseworthy turn out to be closely related.

Thesis Supervisor: Richard Holton
Title: Professor of Philosophy

Acknowledgements

First and foremost I thank my dissertation committee. Richard Holton has been a generous, kind, and patient advisor. I am grateful for all his wisdom, support and his ability to put things in perspective. Alex Byrne has been a source of excellent philosophical guidance and much support – both of the philosophical and the moral kind. Sally Haslanger has provided thoughtful comments and pushed me to be clearer and to think about the big picture. I always left meetings with her inspired. Julia Markovits has been generous with her time way beyond the call of duty. She prompted me to think about many of the questions in this dissertation and her comments and criticism have been invaluable.

I owe a huge debt of gratitude to my “other” dissertation committee: Katia Vavova and Sophie Horowitz. Both me and this dissertation would be much worse off if not for their friendship, enthusiasm, constructive criticism, and epic buddy lunches.

I thank Kenny Walden and Tom Dougherty for comments, helpful discussions, and much encouragement.

I feel incredibly lucky to be part of the MIT philosophy community. I thank Alejandro Perez-Carballo, David Gray, Jennifer Carr, Paolo Santorio, Lauren Ashwell, Nina Emery, Brad Skow, Agustin Rayo, Steve Yablo, Arden Ali, Rebecca Millsop, Rachael Briggs, Dylan Bianchi, Damien Rochford, Brendan Dill, Bernhard Salow, Melissa Schumacher, Ryan Doody, Elisa Mai, Caspar Hare, Jack Marlow-Payne, Heather Logue, Robert Stalnaker, Rae Langton, Roger White, Judy Thomson, Vann McGee, and my excellent classmates Brian Hedden, Dan Greco, Dan Hagen, Miriam Schoenfeld for comments, questions, objections, conversations, jokes, and for making the MIT philosophy department such an inspiring and happy place to be a graduate student.

Bits and pieces of this dissertation were presented at the Cambridge Square Conference 2012, Bellingham Summer Philosophy Conference 2011, Rocky Mountain Ethics Congress 2011, Boston University Ethics Group, MIT WiP Seminar, M.A.T.T.I., University of Wisconsin at Milwaukee, Leeds University, London School of Economics, Cambridge University, and Bern University. I am indebted to their organizers for giving me the opportunity to try out my ideas and I thank their audiences for their thoughtful questions and feedback.

In the course of writing this dissertation, I have benefited from helpful conversations and/or comments on drafts from Daniel Star, Luc Bovens, Ted Hinchman, Josh Schechter, Tyler Doggett, Nomy Arpaly, Ulrike Heuer, Monika Betzler, Justin Broackes, Ned Markosian, David Christiansen, and Miranda Fricker.

Much as I'm grateful to have been surrounded by so many people passionate about philosophy, I'm equally grateful to have encountered quite a few who couldn't care less about it. The MIT Music Department has provided me with a second home and refuge on campus over the years. I particularly thank Bill Cutter, Kerry Deal, and Jean Rife for being excellent and patient teachers.

I thank Maria Giavazzi, Daria Dzyabura, Michaela Bronstein, Chiara Alfano, Henry Swift, Anna Lewis, Erica Bruni, Emma O'Hanlon, and Alex Harris for friendship, support, dinners, and diversions.

I thank my parents, Anna and Peter Sliwa, and my brother, Anton, for their love and support.

I thank Andreas for too many things to list. His unfailing love and support may not have been sufficient for this thesis to be completed. But they were most certainly necessary.

Table of Contents

Abstract	3
Acknowledgements	5
Chapter 1: In Defense of Moral Testimony	9
Chapter 2: Praise without Perfection	33
Chapter 3: Moral Worth and Moral Knowledge	53
Bibliography	83

Chapter 1

In Defense of Moral Testimony

Moral testimony has been getting a bad name in the recent literature.¹ It has been argued that while testimony is a perfectly fine source for nonmoral belief, there's something wrong with basing one's moral beliefs on it. This paper argues that the bad name is undeserved: Moral testimony isn't any more problematic than nonmoral testimony.²

Some people claim that there is something intuitively problematic about deferring to others for one's moral beliefs: there seems to be something valuable about coming to one's moral beliefs by oneself. Hills, argues for example:

Once you have reached maturity as an adult and have the ability to think about moral questions by yourself [...] you have strong reasons to do so, indeed that refusing to do so is unacceptable.³

While children may be in need of moral education and hence should take their parents' word for what's right and wrong, it seems that as adults we shouldn't rely on others for our moral beliefs. Worries about moral testimony are further supported by particular cases, such as this one:

Vegetarian: Eleanor has always enjoyed eating meat but has recently realized that it raises some moral issues. Rather than thinking further about these, however, she talks to 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.⁴

Eleanor's behavior seems disturbing. We can imagine even more disturbing cases:

Suit: Sam is standing at the shore of a lake when he sees a drowning child. He believes that saving the child would be a good thing to do but it would involve ruining his new expensive suit. He cannot decide what to do and there is no one else at the lake, so he decides to call a friend whom he takes to be reliable. His friend tells him that he should save the child, and he believes him and saves the child.

¹ To see that it's been getting a bad name, it's enough to just have a look at the titles. For examples, see Nickel (2001), Hopkins (2007), Hills (2009), Driver (2006), McGrath (2010).

² There has been much recent discussion about the epistemology of nonmoral testimony. See Coady (1996) and Lackey (2008). For the purposes of this paper, I am not committing myself to any particular account of the epistemology of testimony. My argument is just that moral testimony does not differ from nonmoral testimony--whatever the right account of the latter turns out to be.

³ Hills (2009), p. 95.

⁴ Hills (2009), p. 94.

Why are these cases so troubling? Some suggest that what makes these cases peculiar is that they are instances of *moral* testimony. They argue that, were Eleanor to ask her friend about some nonmoral question, this wouldn't be troubling at all. We should hence conclude that there is a general problem about moral testimony. In particular, what explains that the cases are disturbing is a principle like:

NO TESTIMONY: For a mature moral agent, there is something wrong with relying on testimony for one's moral beliefs even if one knows one's source to be reliable and trustworthy.

If something like NO TESTIMONY is right, this has implications for both moral epistemology and meta-ethics. Some take the asymmetry between moral and nonmoral testimony as evidence that what matters in moral epistemology is not moral knowledge but rather moral understanding.⁵ Others suggest that an asymmetry between moral and nonmoral testimony has far-reaching consequences for meta-ethics. In particular, they argue, if moral testimony turns out to be problematic, this can be used as an argument against moral realism in support of non-cognitivist views.

After all, when it comes to deep and unobvious facts about the empirical world, we readily defer to others [...] who are better placed to discover those facts than we are. In such cases, even a very sweeping kind of deference to expert opinion seems appropriate. If, similarly, there is a domain of deep and unobvious moral facts, then it is natural to expect that some of us— intuitively, the “moral experts”—would be better placed to discover these facts than others. In principle then, moral deference should strike us as no more peculiar than deference about scientific matters or geography.^{6,7}

According to moral realist views, moral facts are no different from non-moral facts, the burden is on the moral realist to provide an explanation for any epistemic asymmetry:

⁵ See Hills (2009).

⁶ McGrath (2010), p. 12

⁷ Since I argue that there is no asymmetry between moral and nonmoral testimony, I do not think that there is a special problem for moral realist. But even if there were an asymmetry, this would create a puzzle for moral realists and plausible non-cognitivist views alike. While emotivists have an easy explanation for why moral testimony is troubling, these views are implausible because they cannot account for moral disagreement. More sophisticated views of disagreement, such as Gibbard's norm-expressivism, on the other hand, do not have an explanation of the asymmetry readily at hand. In fact, some of these views explicitly grant that we can rely on others for our moral norms. Gibbard (1990), for example, writes: “When conditions are right and someone else finds a norm independently credible, I must take that as favoring my own accepting the norm.” (p. 180) In general, it seems that any non-cognitivist view that makes room for moral disagreement, does not have an easy explanation for a deep asymmetry between moral and nonmoral testimony.

The goal of this paper is to show that there is no general problem about moral testimony; in fact moral testimony is no more problematic than nonmoral testimony. My strategy is as follows: In Section 1, I give a direct argument that NO TESTIMONY is false. Then, in Section 2, I defend this argument against an objection. In Section 3, I revisit the initial cases that motivated worries about moral testimony and provide a better diagnosis of what goes wrong in them. I show that in these cases the agent's reliance on moral testimony isn't the culprit--rather, the culprit is moral ignorance, controversy, or ulterior motives. Finally, I consider two attempts to spell out and defend NO TESTIMONY. According to the first, moral testimony is epistemically problematic because it doesn't give us moral knowledge. According to the second, moral testimony is morally problematic because it's incompatible with morally worthy actions. I argue that neither of these attempts succeeds. Thus, there is no asymmetry between moral and nonmoral testimony.

1. The Argument from Moral Advice

The initial cases that have motivated NO TESTIMONY are very peculiar, and they suggest that moral testimony is an unusual and exotic phenomenon. But relying on others for our moral beliefs isn't exotic. It's something all of us do by asking for and taking moral advice and it's something we do for good reasons. To get a better picture of the role of moral testimony, it's therefore important to look at a wider range of cases. In this section I will therefore look at some mundane cases in which agents rely on moral advice. I will argue that these cases are intuitively unproblematic and that the agent's reliance on moral advice is a good thing. But, I will argue, there is no difference between relying on moral advice and relying on moral testimony. So any general principle like NO TESTIMONY must be false.

I want to start with some fairly ordinary cases of moral advice.

WEDDING: Tom and Sara are planning a wedding and both of their families have offered to contribute money towards it. Sara's family, who are less wealthy than Tom's, offered a certain sum, which will cover less than half of the expenses. The couple is now wondering whether it would be permissible for them to ask Tom's family (which is wealthier) for a greater contribution. In particular, they worry that it wouldn't be fair of them to ask one set of parents for more. They decide to ask a friend whose judgment they trust.

TRIP: Anna is a journalist who is preparing to go on a reporting trip to a dangerous and conflict-ridden area. She has to tell her family that she will be away but she really doesn't know what to tell them. If she tells them where exactly she's going and why, they will be extremely worried. On the other hand, she worries that by evading the questions she would

be lying. She goes back and forth but cannot decide what the right thing to do is and eventually decides to ask a friend whose judgment she trusts.

FRIENDS: Susan's friends have been playing pranks on a new girl in her class. Susan worries that they might be going too far and that they are bullying the girl. But when she talks to them, they insist that they are just teasing her. Susan doesn't know what to do. If her friends are being bullies, she should step in. But they are her friends and she doesn't want to get them in trouble and she's not quite sure whether what they are doing is bullying. Eventually, she decides to ask a friend for advice.

Intuitively, there isn't anything wrong or disturbing about these cases. In all of them, it seems fine for the agent to rely on their friends' judgments and to take their friends' advice. You might worry that in these cases, the agents aren't really relying on someone else's *moral* judgment. After all, "fairness", "lying" and "bullying" are thick terms, so you might be worried that in asking for advice, the agents aren't really asking for any *moral* information. Now, it does seem true that not all instances of apparent moral advice really do involve a request for moral information. When asking their friend "Would it be okay for us to ask Tom's parents for a greater contribution?" the couple might just be unsure about how Tom's parents would react to the proposal. They might not be so much worried about whether asking more of Tom's parents really *is* fair, but rather whether Tom's parents would *think* that it's fair. And in this case, they aren't relying on their friend's moral judgment but rather on her psychological acuity. Suppose that their friend tells them, "Don't do this, Tom's parents will be upset", and they accept that they shouldn't do it. They might be just relying on their friend for the belief that Tom's parents would be upset and use this nonmoral information to make their own moral judgment.

However, it seems clear that not all requests for moral advice are just requests for psychological information in disguise. We are sometimes uncertain about what the right thing to do is, not because we lack nonmoral information, but because we are genuinely uncertain about the *moral* status of the action or the situation. In particular, we can imagine that Tom and Sara have a very good idea about how Tom's parents would react to their request and they know that they would be happy to contribute a greater part. But even so, Tom and Sara might still be unsure whether it would be *reasonable* to ask them to do so. Similarly, Anna might know that her family will be terribly upset if she doesn't keep quiet about the nature of her trip and they will accuse her of having lied should they find out. Nevertheless, she might be unsure whether they would be *justified* in their accusations. That's because she is unsure about whether what she contemplates

doing really does amount to a lie. And finally, Susan might be aware of the psychological effects that her friends' behavior has on the girl but she might still be unsure if what they're doing really is *bullying* as opposed to just teasing. So, in taking someone's moral advice, we are really relying on their moral judgment. And in the cases I considered, this seems like a good thing.

Why can taking someone's moral advice be a good thing? In relying on someone else's moral judgment, we acknowledge that the other person is in a better epistemic position with respect to the particular moral judgment than we are. And we can have excellent reasons for doing so. Why should we take someone else to be better placed to make a given moral judgment? There are at least two good reasons for doing so: For one, we might be concerned that our own judgment is compromised by bias or self-interest. Secondly, we might think that the other person is just better at making certain moral judgments than we are.

Worries about bias are an important motivation for seeking out moral advice. Figuring out what the right thing to do is often requires us to take an impartial perspective on our action and to bracket morally irrelevant factors in our moral reasoning. But that's not easy. It's hard to know whether or not on a given occasion our reasoning is being influenced by morally irrelevant factors. And then even when we suspect that our reasoning is infected by biases, it's hard to know the extent to which this has happened. So, even if we do realize that our reasoning is biased, we may still not be in a position to fix it. In situations like this, we should rely on moral advice. We should rely on the judgment of someone who isn't subject to our biases. Thus, in *FRIENDS*, Susan might worry that she's not reliable in judging whether the girl is being bullied because it's her friends who are involved. After all, we're less likely to judge our friends harshly and more likely to make up excuses for their behavior. But just knowing this may not help her do better. In fact, if she tries to compensate for the bias, she might even err too much on the side of caution and end up overreacting to even harmless jokes. Thus, asking for advice is the best she can do.

Similarly, in *WEDDING*, Sara and Tom might worry that their own interests can influence their judgment of whether it's fair to ask Tom's family for the greater contribution. After all, if the combined contribution of their parents doesn't amount to the full cost of the wedding, they will either have to pay the rest out of their own (slim) pockets or they will have to change their plans

to cut some costs. So it's in their own interest to think that it would be fair to ask Tom's parents for a greater contribution. Moral judgments often either require us to set our own interests and preferences aside or they require us to weigh self-interested reasons against the interests of others. In such cases, we seek out moral advice because we don't trust our own moral judgment. We take the other person to be in a better epistemic position because they don't share our interests and hence cannot be biased by them.

But we don't just seek out moral advice because we think that our own moral judgment might be compromised. We may rely on someone's moral advice because we believe that they are better at making certain kinds of moral distinctions. We may think that the other person is more sensitive to a given moral norm than we are. This is because seeing whether a given moral norm applies is often a matter of ability. It involves practical knowledge, not just theoretical knowledge. Practical knowledge comes in degrees: For example, some people can just be better at seeing whether some situation is unfair than others. They have finer discriminatory skills and can distinguish situations that aren't obvious to others.⁸

To see what I have in mind, consider a case of social norms:

RUDE: Anna shares the social norm against rudeness and she is trying hard to be polite, but she cannot always tell whether her tone of voice, her behavior or an email she is writing is rude. She worries about this because she doesn't want to be rude. Therefore, whenever, she is uncertain, she relies on her friend's judgement.

It doesn't seem mysterious how Anna could come to believe that her friend is a better judge of rudeness than herself. After all, it's not that Anna lacks the concept of "rude". It's just that her ability to discriminate is rather coarse. Her friend, on the other hand, has finer discriminatory skills. The tone of two emails can look very much the same to Anna, when one of them strikes her friend as clearly more polite than the other. And it's not just that her friend is better at psychological judgments. After all, whether, for example, an email is rude isn't just a matter of whether the recipient will be offended by it. It's a matter of whether the other person is *reasonably* offended by it. Anna, therefore relies on her friend whenever she suspects that finer

⁸ You might worry that in all the unproblematic cases of moral testimony so far, the agent takes on someone else's say-so that a moral norm they already accept applies. Hence, you might worry that I haven't shown that we can acquire a new norm on the basis of moral testimony. I address this worry in the next section. There I show that insofar as accepting a new norm on the basis of testimony seems problematic, it's not testimony that is to blame. Independently of this, I am skeptical that we can draw a sharp distinction between learning how to apply a norm and accepting a norm, in the first place.

discriminatory skills are needed. She trusts her friend's judgments that the tone of an email is rude even if she antecedently didn't think so herself. Anna's predicament doesn't seem unusual. It seems clear that some people are just better at applying social norms than others. But then it doesn't seem implausible that something similar could hold for moral norms, too. Just as Anna could come to believe that her friend is in a better epistemic position to make judgments about rudeness, so we can sometimes come to believe that someone else is better at judging whether a given moral norm applies.

What could make one person more competent with a particular moral norm than others? If, for example, seeing that a situation is unfair is an ability, a form of practical knowledge, then practice and experience matters. Learning to apply a given moral concept may be a lot like learning how to read an X-ray. You have to have seen a whole bunch of them to learn to distinguish the white blotches and patterns that are just artifacts from the ones that are evidence for disease. And the more experience you have, the better you become at distinguishing the tricky cases. Similarly, learning how to distinguish situations in which someone is being treated fairly from the ones where she's being treated unfairly, you may have just had to experience a number of such cases.

The fact that we have different backgrounds means that some of us get to practice applying some moral concepts and making certain moral distinctions more than others. Someone who grew up with many siblings may be more sensitive to considerations of fairness than someone who grew up as a single child. A student from a minority group may be more sensitive to whether a given remark is racist than her white peers. Thus, even if they believe that racism is wrong, her peers may still want to defer to her whether some particular joke is racist or merely distasteful. Similarly, one may want to defer to one's female friends whether some particular remark is sexist, even if one is fully committed to combat sexism.⁹ Of course, there is no easy route from background to competence with a given moral norm. The mere fact that someone grew up in a large family, doesn't guarantee that they will be better judges of fairness just as being a woman doesn't guarantee that someone will be a good judge of sexism--she might see sexist remarks

⁹ The central example in Jones (1999) is also of this kind. In the example, Jones argues, Peter ought to defer to his (female) roommates' judgments of sexism. Jones uses this example to argue for a narrower conclusion than I do. I argue that *in general* there isn't anything wrong with relying on moral testimony, not just that in some particular instances there isn't.

where there are none. But it still seems plausible that *some* backgrounds can make a person more reliable in applying some particular moral norm.

If making some moral distinctions is in fact a matter of practical knowledge, this gives us a further reason to seek out advice from someone whom we regard as better at them than we are. Not only will it help us to do the right thing but it might help us become better at making the distinctions ourselves. If you're learning how to read an X-ray, you do well in asking your more experienced supervisor for advice in cases in which you are unsure. Similarly, in relying on her friend, Anna might actually get better at recognizing rudeness herself. And trusting one's female friend's judgments about sexist remarks, might help one learn how to make that distinction yourself.

I have argued that moral advice is a good thing for two reasons: It's morally good because it helps us do the right thing in cases in which we might otherwise fail to do so. And it's epistemically good because it allows us to take advantage of our peers who may be better epistemically placed than us to make certain moral distinction and to come to the right moral conclusion. In short, we ask for and rely on moral advice because we're in many respects creatures who make mistakes, who get distracted, who are susceptible to biases, who have limited abilities of discrimination in some areas and who are, moreover, well aware of all that. Moral advice allows us to do the right thing despite all these limitations by tapping into the cognitive resources of our peers.

Note that we rely on nonmoral advice for exactly the same reasons. Imagine a doctor asking her colleague for advice for help in making a diagnosis on the basis of an X-ray. She might be requesting help because of concerns about bias: The patient is her friend and she worries that this might influence her judgment, since she really hopes that her friend is healthy. Or she might be worried about not having enough experience to interpret correctly the result--she just hasn't had much practice looking at this particular kind of picture. Or she might just think that she's overworked and sleep deprived and therefore doesn't quite trust her own medical judgments. In taking medical advice from her colleague, the doctor relies on testimony. She accepts a belief on her colleague's say-so because she takes her colleague to be more reliable than herself. Just in the same way, the instances of of moral advice that I discussed above involve moral testimony: In

taking someone's moral advice, you often accept a moral belief on the basis of their say-so. (I will defend this claim further in the next section.) It seems clear that in taking medical advice, the doctor isn't doing anything wrong. Similarly, it seems clear that in the situations I described above, there isn't anything wrong with taking your friend's moral advice. But if there isn't anything wrong with taking moral advice and there is no difference between taking moral advice and accepting moral testimony, then NO TESTIMONY must be false. There isn't anything particularly problematic about relying on others for your moral belief.

2. Is Moral Advice just Moral Testimony?

I have argued that there are many instances of moral testimony in the form of moral advice that are unproblematic. Therefore, NO TESTIMONY must be false: there is no general problem about moral testimony just as there is no general problem about nonmoral testimony. In this section I respond to an objection to my argument. According to Hills, moral advice and moral testimony are distinct and importantly different. Hence, the fact that moral advice is unproblematic doesn't establish that moral testimony is.

How, according to Hills, does moral testimony differ from moral advice? Hills argues that the two differ in that, in taking moral testimony,

...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.

According to Hills, once you have determined that the speaker is reliable

...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.

In contrast, in taking moral advice you don't just "believe what is said to you." Rather,

you subject [the testimony] 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.¹⁰

But this distinction between moral testimony and moral advice is not plausible and it doesn't track the distinction between problematic and unproblematic cases. For one, there are cases of moral testimony which are unproblematic but which do not involve treating the testimony

¹⁰ Hills (2009), p. 122-123.

merely as a guide to one's own reflections. I gave several examples of such cases above. In these examples, the agents rely on moral advice because they want to do the right thing, but they either worry about the reliability of their own moral reasoning or they think that someone else is better at making certain moral judgments than they are. In both cases, the testimony plays a much more substantial role than just "guiding reflection", yet the agent doesn't seem to be doing anything wrong in relying on it. Would it be better if the agents did engage in further critical deliberation? No, in fact, it seems that in cases like these, further deliberation after receiving the testimony may be exactly what the agent *shouldn't* do. After all, if the agent is worried about her moral judgment being biased, there is no guarantee that her further deliberation won't be biased or overly impressed with her self-interested reasons as well. And once she starts reflecting further, she opens the way to succumbing to temptation and rationalizing the testimony away. In these situations further reflection may just be an invitation to fall prey to weakness of will.¹¹ What about the case in which the agent takes moral advice because she thinks that her friend is better at making the relevant moral distinctions? In these cases, too it's not clear why critical deliberation would be called for because it's not clear how critical deliberation could help. Moral advice is so useful precisely because it's a means to put an end to one's reflections about what the right thing to do is.

Thus, there are plenty of instances of moral advice which don't just involve treating moral testimony as a guide to one's own reflection and in which nevertheless the agent doesn't do anything wrong. On the other hand, however, testimony, both moral and nonmoral, does often involve more critical reflection than Hills' account seems to suggest. In fact, without further reflection, reliance on testimony might not be epistemically rational. To see why, it's helpful to consider an extreme, nonmoral case. Suppose I go to the doctor, whom I take to be reliable and trustworthy, to ask him what to take for my headache. The doctor tells me that my headache could be cured by taking a generous spoonful of cyanide. Now, even if prior to his testimony, I regarded the doctor as reliable and trustworthy, I would be epistemically irrational (and most likely dead) if I "simply believed what was said" to me. That's because even reliable testimony is in general just one piece of evidence that I have. In this case, I have other evidence about headaches and the likely effects of cyanide. What I need to do is to weigh the evidence from testimony against all the other non-testimonial evidence that I have. That requires critical

¹¹ See Holton (2010).

reflection about the testimony itself. Even in cases in which a speaker's testimony is the only evidence I have about some question, epistemic rationality may still require me to think about the plausibility of what I have been told. After all, even someone who's usually reliable and trustworthy can be tired or drugged or joking on that particular occasion. Asking "What makes you think so?" is an easy way to check whether the speaker has really thought the problem through, whether she really has got your question right and to rule out that she's not just trying to get rid of us, or joking or drunk. That's why we often don't just simply believe as we're told without any further questions and that's why scenarios in which the agent doesn't do so may seem intuitively odd.¹²

Trying to draw a distinction between advice and testimony by taking one but not the other to involve critical reflection seems wrong-headed in both moral and non-moral cases. In testimony there is sometimes quite a bit of critical deliberation, as when you convince yourself that the doctor really knows what he's doing when he is suggesting you take cyanide for your headache. Taking advice sometimes involves very little, for example when you don't trust your own judgment on the particular question. And there are plenty of cases in between, depending on whether you ask for advice because you need to put an end to your reflection and come to a decision, whether you are asking for advice because you want some help with thinking through the problem, and depending on just how much help you want. Moral advice and moral testimony aren't fundamentally different. Rather, moral advice just is a subclass of moral testimony. It's testimony about practical questions: questions about whether we ought to do something, whether it would be a good thing for us to do it, or what a good way of doing something is. Since many moral questions are practical questions, questions about what we should do, much of moral testimony takes the form of moral advice.¹³

3. Ignorance, Disagreement and Ulterior Motives

I have argued that there are plenty of unproblematic cases of moral testimony, so we should reject the principle NO TESTIMONY. But what about the initial cases, such as *Suit* and *Vegetarian*, that motivated NO TESTIMONY? I agree that these cases are problematic. But, I will argue in this

¹² See, Christensen, David (2007). Epistemology of disagreement: The good news. *Philosophical Review* 116 (2): 187-217.

¹³ For example, advice is often presented in the form: "Here's what *I* think..." In giving advice, the speaker hence often doesn't speak as authoritatively as when she is testifying.

section, they aren't problematic because of any general problem with moral testimony. I will isolate three things that can go wrong when an agent asks for moral testimony. I will argue that these factors can explain our intuitive unease about the problematic scenarios. But they don't show that there is anything wrong with moral testimony per se. In particular, I will argue, these factors can make cases of nonmoral testimony just as problematic.

Moral Ignorance

In general, we ask for moral advice because we don't know what the right thing to do is. Testimony is primarily a means for sharing information and resolving ignorance. The request for moral testimony is then evidence for an agent's moral ignorance. Moral ignorance is complicated: there are many reasons why an agent can fail to know what the right thing to do is. I will argue that it's the moral ignorance of an agent that drives our intuitions in some of the more troubling cases of moral testimony. What our intuitions are latching onto in these cases isn't that the agent is resolving her moral ignorance by relying on testimony, but rather that she is morally ignorant in the first place. Some instances of moral ignorance are just problematic in and of themselves. But, I argue that once we're clear on the fact that it's the ignorance that's problematic, we will see that there isn't any further problem about resolving the ignorance by testimony. On the contrary, in some of the most problematic cases of moral ignorance that's exactly what the agent should do.

Recall one of the initial cases:

Suit: Sam is standing at the shore of a lake when he sees a child beginning to drown. He believes that saving the child would be a good thing to do but it would involve ruining his new expensive suit. He cannot decide what to do and there is no one else at the lake, so he decides to call a friend whom he takes to be reliable. His friend tells him that he should save the child, and he believes him and saves the child.

This request for testimony doesn't just seem odd--it seems outright bizarre. There must be something very wrong with Sam if he cannot see that saving the life of a child outweighs saving one's expensive suit. Or consider the following:

Ron, the extremist, would like to kill Tamara because he is angry at her but he isn't sure whether that's morally permissible. He decides to ask his Rabbi what the right thing to do is. The Rabbi tells him that he mustn't kill Tamara and Ron believes him.¹⁴

¹⁴ A similar case appears in Hills (2009).

Again, the fact that one mustn't kill someone just because one is angry with them is an obvious moral truth. It's not exactly moral rocket science. To be ignorant of such basic moral truths, one's sensitivity to moral reasons must be seriously compromised. At the same time however, these agents aren't outright *bad*. They aren't psychopaths who have no concern for morality whatsoever. They *are* motivated to do the right thing, they're just very bad at figuring out what the right thing is. But again, the fact that they are bad at figuring this out doesn't escape them. They do worry about getting it wrong and it's for that reason that they ask for moral advice.

What's troubling then about these cases is the agents' ignorance of basic moral truths. But are they in addition doing something wrong by relying on moral testimony for their beliefs? This just doesn't seem plausible at all. For one, it's just not clear what a better alternative would be. After all, left to their own devices, these agents are likely to come to the wrong conclusions and make terrible moral mistakes. So relying on testimony is exactly what they should do. We might worry that someone who is so insensitive to moral facts might also be no good at recognizing people who are reliable in their moral judgments. But even if they aren't terribly good at assessing the reliability of other people's judgment, they can still rely on moral testimony since almost anyone is in a better position to make moral judgments than they are. Their situation is rather like someone's who suffers from severe colorblindness and must therefore rely on others for accurate color judgments. Even though such a person isn't well posed to evaluate other people's reliability about color judgments, her reliance on testimony isn't problematic. That's because almost anyone's color vision is better than hers.

I argued that what makes some of these extreme cases of testimony problematic isn't the testimony but rather the ignorance that motivates the request for it. And it seems that this isn't any special feature of moral testimony, nor is it a special feature of moral ignorance. Nonmoral ignorance can be problematic and when it is it can make cases of nonmoral testimony look intuitively odd, too. Consider the following:

Mark, a sexist, believes that women are less intelligent than men. He comes across plenty of conclusive evidence that this is not so--in his college classes there are plenty of women who are smart, the valedictorian in his year is a woman, etc. The conclusive evidence doesn't manage to convince him of the contrary but he does undermine his confidence a bit and he decides to consult a (male) friend whom he takes to be reliable and trustworthy. His friend tells him that women aren't less intelligent than men and Mark believes him.

Or:

Susan is wondering whether sense perception is a reliable source of belief. She decides to ask a friend whose judgement she trusts. Her friend tells her that sense perception is reliable, so she believes him.

These cases look odd. It's not that there is a general principle prohibiting one to ask how women perform on IQ tests. But in this case Mark shouldn't have to ask. He doesn't know that women aren't any less intelligent than men because he's a sexist. So even though his uncertainty is about an empirical claim, his ignorance is problematic. Were he not in the grip of his sexism, the facts would be immediately obvious to him. His sexism makes him insensitive to the evidence.

Susan's case also seems very strange: It's just odd that a reasonable person would be genuinely uncertain about whether their sense perception is a reliable source of belief. Susan has clearly overdosed on some of the skepticism literature and has lost touch with reality. Now, even though these cases are odd, they wouldn't lead us to conclude that there is something problematic about nonmoral testimony. But then we shouldn't be tempted to draw this conclusion in the moral case either.

Controversy

Consider the other initial case that motivated NO TESTIMONY:

Eleanor has always enjoyed eating meat but has recently realized that it raises some moral issues. Rather than thinking further about these, however, she talks to 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.

Eleanor is uncertain whether eating meat is morally permissible. But her ignorance doesn't seem to be what's problematic in this case. After all, the question of whether it is permissible to eat meat isn't quite as straightforward as the question whether one should save a drowning child or whether one may kill someone with whom one is angry. It is a question that is subject to much controversy. And it's a complex issue: It involves questions about the moral status and rights of animals, about farming practices, about the environmental impact of animal farming, and the impact of your eating habits. A lot of intelligent people, who are concerned about doing the right thing have thought hard about the issue and nevertheless come to different conclusions. All

this makes it hard to see why there should be something wrong with someone who's aware that there are moral issues about eating meat and just isn't sure how to weight those issues. And presumably, that's the situation that Eleanor finds herself in: She realizes that there are moral problems with eating meat but she's unsure what to conclude. The ignorance is a consequence of the question being hard rather than of some serious defect on Eleanor's part. Insofar as this scenario is problematic, moral ignorance doesn't seem to be the culprit.

The problem in this case has to do with Eleanor's resolving her moral ignorance by testimony. In particular, it's precisely the factors that make Eleanor's ignorance unproblematic that rule out testimony as a means to resolving her ignorance. The question whether it's wrong to eat meat is a controversial question: it's something even intelligent, well-informed and thoughtful people disagree on. But given that there is so much disagreement about the issue at hand, testimony just isn't a reliable source of moral belief. That's because Eleanor cannot expect an arbitrary friend, even if she is generally reliable, to be reliable *on this particular issue*. The friend whom Eleanor asks happens to believe that eating meat is wrong. But there are plenty of reasonable, trustworthy and generally reliable people out there who believe otherwise and had Eleanor happened to ask any of those, she would have ended up with the opposite belief.¹⁵

Again, this isn't specific to moral testimony. Take the following case:

You have bought what you took to be a real Monet for what you thought was a very good price but now that you have hung the picture up in your living room, you are wracked by doubts about its authenticity. You decide to consult the experts and it turns out that your painting has been subject to a lot of controversy. Some distinguished experts insist that it's indeed a real Monet while others argue that it's a fraud. The debate has been raging for some time and both sides have pointed to evidence to support their view.

Can you rely on testimony by one of the experts to resolve your uncertainty about whether your painting is real or a fake? It seems obvious that you can't. This is a controversial question. Much as you'd like to sleep in peace, you can't just go with one of the sides. This brings us back to the issue of expertise that I discussed above. Clearly, in this case, the fact that the authenticity of your painting is controversial and you're not in a position to ascertain which of the purported experts has it right doesn't establish that there are no art experts or that you are never justified in

¹⁵ This also explains why it would be epistemically problematic to accept many moral norms on the basis of moral testimony. If, for example, you became a consequentialist on the basis of moral testimony, your belief would not be justified.

trusting one. It's just that you cannot trust their testimony on controversial questions because you have no means to identify who is right in *these* cases. Similarly, the problem with Eleanor's reliance on moral testimony about Vegetarianism isn't a general problem about ascertaining moral reliability. Rather, the problem is that Eleanor isn't in a position to identify who is right *about the particular question* of whether she ought to eat meat. She has no means of figuring out who got it right not because it's a moral question but because it's a controversial question. It doesn't imply that moral testimony is problematic. It's just a consequence of the principle that testimony is not a reliable source of belief about questions that are controversial--a principle that rules out *some* cases of moral and nonmoral testimony alike.

Ulterior Motives

Finally, there is a third factor that can explain our intuitive discomfort with some requests for moral testimony. In all the cases of testimony that I have focused on so far, the agent who asks for advice is motivated by her desire to do the right thing and her uncertainty about what this is. But not all requests of testimony are motivated by the desire to do what's right. Sometimes there are ulterior motives at work. Consider the following case:

Whenever she has to make any hard moral decisions, Susan always asks her mother for advice. When it then turns out that she has made the wrong choice, Susan immediately blames it on her mother, saying that she only acted this way because her mother told her to.

Here, Susan isn't requesting moral advice because she wants to do the right thing and she's uncertain what that is. In fact, even though it may look that way, Susan isn't really asking for moral advice at all. She's not relying on her mother to resolve her moral uncertainty. Rather, she is asking her mother to make the decision she faces *for* her. And she's not doing so out of concern for doing what's right. She just wants to blame someone else in case something goes wrong. She is trying to avoid responsibility for her actions.

Or consider again *Vegetarian*:

Eleanor has always enjoyed eating meat but has recently realized that it raises some moral issues. Rather than thinking further about these, however, she talks to 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. (Hills)

And suppose we fill in some more background: Eleanor is an extremely self-conscious teenager who is trying as hard as she can to fit in with her peers. She is mortified at the mere thought of saying or doing anything that would make her stand out. With this additional information, it becomes plausible that what motivates Eleanor to ask her friend about whether or not she should be a vegetarian is her desire to conform to her peers rather than her concern for doing the right thing. It may look like she is asking for moral advice, but she isn't really: she's interested in whether her friends *think* it's wrong to eat meat not whether it actually *is* wrong. Just like Susan in the case before, Eleanor is asking for someone else to make a decision for her and she's motivated by considerations other than concern for doing what's right. What Susan and Eleanor are doing is clearly problematic. It's problematic because they are relying on other people to make their decisions for them. But the cases aren't problematic because they involve moral testimony. In fact it's easy to imagine cases which involve nonmoral testimony:

John is working in a big lab. He has a big ego and he is very worried about making a mistake and displeasing his supervisor. He therefore doesn't even attempt to interpret the experimental data by himself. Rather he completely relies on his lab partner. When a mistake happens he immediately blames him.

John's deference to his lab partner is no less troublesome than Susan's or Eleanor's even though it isn't deference about a moral proposition. It's problematic because he is relying on his lab partner for making decisions for him, so that he doesn't have to take responsibility for his actions.

Consider yet a different case:

Maria is very timid and has absolutely no self-confidence. She thinks that she can never do anything right. Whenever something she does, turns out badly, she chastises herself and takes this as evidence that she's stupid and no good. When it does turn out well, she thinks that she just got lucky. One day, she decides to just give up. Henceforth, she relies on her husband for all her moral decisions and she completely defers to him.

This case, too, is very troubling. Maria is completely reliant on her husband--not just for moral testimony but for her decisions. She isn't taking her husband's advice; rather, she is using him to decide *for* her what to do. But unlike in the previous cases she isn't doing this out of a desire to avoid the consequences of her action but because she lacks self-worth. Maria's autonomy is compromised. She seems to have given up as an agent. So again, what's problematic about this case isn't moral testimony. In fact, the case seems equally problematic, if we imagine Maria deferring to her husband on all kinds of nonmoral matters: On how to dress, what to eat, how to

stay healthy. These cases then may look like instances of testimony but they shouldn't really be classified as such--unlike in testimony, the agents here are using other people to make their decisions for them rather than relying on them for belief, so they can decide on their own.

4. Defending NO TESTIMONY: Moral Expertise and Moral Worth

In this section, I consider two attempts to defend NO TESTIMONY. The first defends an epistemic version of NO TESTIMONY, according to which moral testimony is problematic because it cannot give us moral knowledge. According to the second attempt, moral testimony is morally problematic because relying on testimony is incompatible with an agent's actions having full moral worth.

4.1 Moral Expertise

Some people have argued that there are epistemic reasons to worry about moral testimony. In particular, some argue that in order to be in a position to rely on testimony, one must be in a position to identify experts. But, according to this view, moral expertise is problematic: even if moral experts exist, it's hard to see how we could identify them. Why think that there is something especially problematic about moral expertise? McGrath suggests that we identify experts by asking them controversial questions and checking whether they give the right answer. However, we can't do this in the moral case, since we have no independent access to moral facts:

By contrast [*with nonmoral judgment*], there seems to be no analogous way to calibrate the accuracy or reliability of someone's moral judgment, because one lacks the relevant kind of independent access to the moral facts. If one attempted to rank others with respect to the accuracy of their moral judgment by checking how often they answered controversial moral questions correctly, it seems as though one could do so only by engaging in first-order moral reasoning and deliberation of one's own. It is thus unsurprising that clear and unequivocal evidence that someone possess unusually reliable moral judgment is hard to come by. Even if there are genuine moral experts, locating particular individuals within the space of moral expertise is undoubtedly a precarious business.¹⁶

I will argue that the proposal is not a good method for identifying any kind of experts, moral *or* nonmoral. I will then argue that the absence of an "independent check" cannot explain why it's supposed to be problematic to rely on moral testimony. Moreover, we don't need an

¹⁶ In her (2009) McGrath argues that our inability to identify experts accounts for the asymmetry between moral and nonmoral testimony. Later, in her (2011), she argues that it's only part of what makes moral testimony is especially problematic -- the full explanation also involves the problem of morally worthy actions on the basis of testimony.

“independent check” in order to rely on testimony. If we did, this would impugn much of nonmoral testimony, too. Finally, I will show that if expertise is required for testimony, it’s only a very weak kind of expertise. To be justified in taking someone’s word on some question, all you need--at most--is to have reason to believe that the question is reliable on *this particular question*. I argue that we easily can have such reasons.

Asking controversial questions is not a plausible proposal for identifying experts of any kind. Suppose you were trying to figure out whether you should trust your doctor about which headache medication to take. According to McGrath, what you have to do, is to establish your doctor’s medical expertise. And the way you go about establish your doctor’s medical expertise, on McGrath’s view, is by asking her a number of controversial medical question--so maybe you’d have to ask your doctor about the cure for cancer or whether allergies are caused by environmental pollution--and see whether she answers them correctly. But, of course this isn’t a test that you are in a position to carry out and, importantly, this is not just for the lack of medical background and training. In fact, if what it takes to identify a medical expert is to see whether or not she answers controversial medical questions correctly, then it seems that no one is in a position to identify medical expertise. That’s because it’s in the nature of controversial medical questions that even people with the requisite training and experience don’t know what the correct answers to them are. If it were possible to just check who answered a controversial question correctly, it wouldn’t be a controversial question.

The requirement of an “independent check” in order to determine expertise is problematic. Recall that what we wanted to explain is why it seems that an agent should come to her own moral conclusions, rather than rely on others for her moral beliefs. But if what makes moral testimony epistemically problematic is the lack of an “independent check,” then it’s hard to see why it would be any better for an agent to rely on her own judgment rather than on that of her friend’s. After all, the lack of an “independent check” doesn’t just preclude her from assessing other people’s reliability. It also makes it impossible to assess her own reliability--she is no more in a position to check whether she herself got it right than she is in a position to check whether someone else got it right. Hence, relying on others or relying on herself is epistemically on a par. Thus the fact that we don’t have an “independent check” for moral truths, cannot explain why relying on moral testimony, rather than on one’s own judgment, is epistemically suspect.

More importantly, it's simply implausible that we should need an "independent check" in order to be able to rely on testimony. I can rely on other's testimony about what they saw and heard. But I am not in a position to assess the reliability of their perceptual faculties without relying on *any* deliverances of my own senses, just as I cannot assess someone's moral judgments without relying on moral reasoning of my own. Similarly, I can rely on someone's mathematical testimony. But I don't have "independent access" to mathematical facts and the only way I can assess your reliability is by relying on my own mathematical reasoning or the mathematical reasoning of others' whom I trust.

Thus, in order to be epistemically justified in relying on someone's testimony in some domain, you need neither establish that they can correctly answer controversial questions in that domain, nor do you need to have "independent access" to facts in this domain. You need to have reason to believe that the other person is reliable on *this particular issue at hand*. It's not hard to see how you could come to have such a justified belief. As I have shown earlier, you could easily have good reason to believe that the other person is more reliable than you are with respect to some moral question, either because you think that your own judgment is impaired in this particular case or because you have reason to think that the other person is better at making the particular moral distinction at hand. Maybe you have seen the person make good moral judgments before and you know that they have thought about the issue at hand. Or maybe they have been recommended to you as a good advisor by someone whose judgment you trust. Maybe you have asked them some related moral questions and seen that they give reasonable answers. None of this is very different from how you are justified in relying on others' for nonmoral testimony.¹⁷

4.2 Moral Worth

I want to consider one final attempt to defend NO TESTIMONY: moral testimony is morally problematic because it's incompatible with morally worthy actions. On recent accounts of moral worth, an action is morally worthy only if the agent performs it for the reasons that make the action right.¹⁸ Hills and McGrath, have both argued that if an agent relies on testimony then

¹⁷ See also Anscombe (1981), pp. 46-47.

¹⁸ In particular, it's not enough that the agent performs the action because she believes that this is the right action. For a defense, see Arpaly (2003), and Markovits (2010).

her action will not have moral worth even if she does the right thing. They argue that in order to be motivated by right-making reasons, the agents needs to have moral understanding. This requires the agent to have beliefs about *the reasons* which make the action right, not just the belief that the action *is* right.¹⁹ As, McGrath argues:

if an agent Φ 's because of her belief that Φ -ing is the right thing to do, but she does not understand why Φ -ing is the right thing to do, this detracts from the status of her *action*.²⁰

The assumption that moral understanding is necessary for morally worthy actions is problematic, though I will not argue against it here.²¹ But I will argue that moral testimony is not an obstacle to morally worthy actions, even if morally worthy actions do require moral understanding. For one, moral testimony can help you acquire moral understanding. Secondly, even in cases in which testimony itself is not sufficient to acquire moral understanding, moral testimony is not in the way of your actions' having moral worth.

What is moral understanding? Hills argues that moral understanding involves a set of abilities: the ability to give and follow explanations and articulate reasons for an action's moral valence.²² According to Hills, the fact that understanding involves these abilities, and that moral understanding is valuable gives us instrumental reason not to rely on moral testimony:

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.²³

However, it's puzzling, why we couldn't learn to explain why an action is wrong on the basis of moral testimony. After all, I can tell you that what you are about to do is wrong because it's cruel or unkind and thus, on the basis of my testimony, you can come to believe both *that* your action is wrong and *why* it's wrong. Hence, it's hard to see why, given the importance of moral understanding, we could have "strong reasons" to avoid trusting moral testimony. Even if in some cases testimony may not be sufficient in order to acquire understanding, it doesn't follow that we have reasons to avoid it. In fact, testimony may sometimes be necessary in order to achieve moral

¹⁹ For another defense of this view, see Nickel (2001).

²⁰ McGrath (2010), p. 38-39.

²¹ For objections, see Markovits (2010).

²² See Hills (2009), p. 102-103

²³ Hills (2009), p. 98

understanding. This is no different in the moral case than in nonmoral cases. We value not just moral but also, for example, scientific understanding. And while testimony may not always be sufficient to acquire understanding of physics, testimony is nevertheless crucial--for example through textbooks and lectures.²⁴

But suppose you act on moral testimony in a case in which moral testimony is not sufficient for you to acquire moral understanding. Does your action then have moral worth? Note that moral testimony is a means to resolving moral uncertainty in situations. When I am uncertain about whether *p* is the right thing to do, I have moral reason to resolve my uncertainty in a way that will make it most likely that I get it right. If I reasonably believe that someone else is in a better epistemic position to get the right answer than I am, then I have moral reason to ask that person for moral advice. Thus, an agent who seeks out and trusts reliable testimony because she wants to get it right *is* doing the right thing (namely, resolving her uncertainty) *for the right reasons* (she needs to resolve her uncertainty to do the right thing) and she has *moral understanding* (she reasonably believes that someone else is in a better epistemic position to find out what the right thing to do is). Her seeking out and trusting moral testimony is thereby a morally worthy action. If she didn't seek out the testimony or didn't take it, even though she believed it to be reliable, she would be blameworthy for it. And she would be even more blameworthy if she avoided or failed to trust moral testimony that she regards as reliable because she wanted her action to be morally worthy.

So, asking and taking moral advice is an action that can be morally worthy and for which an agent can deserve praise or blame. But McGrath's and Hills' concern was with the action that results from moral testimony. Even if asking for and accepting moral testimony that one regards as reliable is morally worthy, is the action that results from it also *maximally* morally worthy? If Sam relies on his friend's testimony for his belief that he ought to jump into the lake and save the child, is his saving the child a morally worthy action?²⁵

²⁴ See also Annas (2003).

²⁵ See Markovits, (forthcoming) for a defense that acting on moral testimony is morally worthy because testimony is a right-making reason.

Even if an action that is based on moral testimony isn't *maximally* morally worthy, this does not show that the obstacle to moral worth is moral testimony.²⁶ Rather, given the agent's moral ignorance, performing a maximally worthy action simply may not be one of her options. Consider again Sam's case: Suppose Sam doesn't call his friend and hence doesn't ask for moral testimony. Instead, he relies on his own reasoning even though he knows that his friend is more likely to get to the right answer than he is. But suppose that he does happen to rightly conclude that he needs to save the child. It's not clear that in this case his saving the child really is morally praiseworthy--after all, he took a risk in relying on his own deliberation. And even if it is morally praiseworthy, at the same time, he deserves blame for not seeking out moral testimony. After all, he had reason to think that his friend was more likely to find out what the right thing is than he was. In this case, if there is an obstacle to his action's having maximum moral worth, this obstacle is Sam's ignorance rather than his reliance on moral testimony.

Worrying about whether an agent's action deserves maximum moral credit seems a little odd. It seems equally strange to worry whether Sam's saving the child is morally worthy as it does to worry if someone got full moral credit for saving the child, if she called out to a friend to jump in because she was not a confident swimmer. Ultimately, what matters is that we do the right thing. Given that we are cognitively limited, asking for help when you take someone to be better at figuring out what the right thing to do is, is often the best we can do. Of course, a perfect moral agent who knows all moral and nonmoral facts, wouldn't have to rely on moral testimony. In fact, she wouldn't have to rely on *any* kind of testimony. She would always be in a position to do the right thing for the right reasons. But it would be very silly of us to try to do as the ideal moral agent does.

5. Conclusion

I argued that moral testimony is no more problematic than nonmoral testimony. This shouldn't be surprising: Nonmoral testimony is so valuable because we are fallible creatures with limited cognitive resources. Moral testimony is valuable for precisely the same reasons. And we rely on moral testimony just as we rely on nonmoral testimony. This doesn't mean that there isn't anything wrong with the initial cases that motivated the principle NO TESTIMONY. These cases are problematic, but I argued that they aren't problematic in virtue of being cases of *moral*

²⁶ I argue that an action based on moral testimony can have maximum moral worth, elsewhere.

testimony. I suggested a better explanation what goes wrong in them and I argued that cases of nonmoral testimony can be problematic for exactly the same reason. Finally, I argued that moral testimony is compatible with morally worthy action. In fact, I suggested, asking for moral testimony can sometimes be the best we can do.

Chapter 2

Praise without Perfection

Escorted by a score of photographers from various gossip magazines, a celebrity delivers a check to an African village. We're not too impressed with her action: there's no question that charitable giving is the right thing to do, particularly if you are as wealthy as the celebrity. But still, we don't consider her action morally praiseworthy. A young doctor hears about the terrible health conditions in a poor country in the aftermath of a natural disaster. Rather than relaxing on the beach, she spends her vacation volunteering there. Unlike the celebrity, she does seem morally praiseworthy for her action.

What distinguishes those right actions that merit moral praise from those that do not? What makes the actions of the young doctor different from those of the celebrity? This is what an account of moral worth tries to answer. To say that an action has moral worth just is to say that the agent is morally praiseworthy for performing it. The following is a tempting thought: to have moral worth an action needs to be done *for the right reasons*. Unlike the doctor, the celebrity does not act for the right reasons; she is just motivated by self-interest. This tempting thought has been popular in recent literature. Arpaly, for example, takes it to express "an important truth about moral worth":

For an agent to be morally praiseworthy for doing the right thing is for her to have done the right thing for the relevant moral reasons—that is, the reasons for which she acts are identical to the reasons for which the action is right.²⁷

A popular strategy in recent literature has been to turn the slogan into an account of moral worth. This requires giving answers to two questions: What are the *right reasons*? And, given an answer to that question: what is it for an agent to act *for* them?²⁸

The aim of this paper is to examine how promising this approach to moral worth is. To do so, I consider three of what I take to be the most promising proposals that try to give an account of moral worth in terms of "acting for the right reasons"—one attributed to Kant, one by Arpaly, and one by Markovits. These proposals differ in what they take to be "right reasons" and what it is

²⁷ Arpaly [2003], p. 72. Similar theses are accepted by Markovits [2007] and Stratton-Lake [2003]. Markovits bases her account on the "Coincident Reason Thesis", according to which "my action is morally worthy if and only if my motivating reasons for acting coincide with the reasons morally justifying the action" (p. 205). Stratton-Lake defends the "Symmetry Thesis", according to which "the reason why a morally good person does what she should, and the reason why she should do that act are, under favourable conditions, the same." (p. 3)

²⁸ For example, Arpaly [2003], Stratton-Lake [2000], Korsgaard [2008], and Markovits [2007].

to act for them. I will argue that none of them are satisfactory; they all fail to accommodate some cases of intuitively morally praiseworthy actions. The problem, I will argue, is to reconcile the requirement of “acting for the right reasons” with the fact that agents who fall short of moral perfection in various ways can be morally praiseworthy for their actions. Giving an account of “right reasons” that is compatible with these imperfections turns out to be a challenge that none of the accounts I consider manages to meet. This, of course, does not conclusively establish that it cannot be met. But it does motivate looking for an alternative approach. I develop such an approach in *Chapter 3*.

The paper proceeds as follows: *Section One* considers a view that has been attributed to Kant. According to this view acting on right reasons requires an agent to perform an action because it’s the right thing to do. *Section Two* presents a counterexample to this view: the case of Huck Finn. This counterexample motivates Arpaly’s view, which I discuss in *Section Three*. According to Arpaly, moral worth requires us to act from concern for right-making reasons. In *Section Four*, I argue that this view, too, is subject to counterexamples. In particular, it cannot accommodate cases in which an agent relies on moral testimony. In *Section Five* I argue that we have a general challenge for identifying the “right reasons” relevant for moral worth: moral worth must be compatible with an agent’s falling short of moral perfection. I consider a view that tries to meet this challenge in *Section Six*. This view, defended by Markovits, proposes an alternative account of right-making reasons in terms of evidence. *Section Seven* argues that this view, too, is subject to counterexamples.

1. Concern for Morality

In the *Groundwork*, Kant argues that an action has moral worth only when it’s performed from the motive of duty:

in the case of what is to be morally good it is not enough that it *conform* with the moral law but it must also be done *for the sake of the law*. (4:390)²⁹

Kant contrasts “acting for the sake of the law” with two other ways of being motivated. To act for the sake of the law is different from being motivated by one’s self-interest. Hence, the actions of a shopkeeper who treats his customers honestly do not generally have moral worth. This is because shopkeepers want to do what’s good for their profit when making business decisions. So when

²⁹ Parenthetical citations from Kant’s work refer to the Royal Prussian Academy of Sciences edition of Kant’s *Gesammelte Schriften*, with volume and page numbers separated by a colon.

the shopkeeper decides to price his goods fairly this isn't because he wants what's right. It's because being honest in one's dealings with customers happens to be, as a matter of fact, good for business. In this particular instance business interests and morality recommend the same course of action. But they could come apart. And if they were to come apart the shopkeeper, insofar as it's his desire for profit that moves him, would be led to do the wrong thing and cheat his customers.

Kant also contrasts actions performed from the motive of duty with actions performed from "inclination". Consider for example the sympathetic shopkeeper who treats a customer honestly from sympathy. According to Kant, sympathy is a good thing; it often makes acting from the moral law more pleasant. But it's not sufficient to confer moral worth on an action. The agent needs to be moved by her concern for morality; not by, for example, her desire for the well-being of the customer. The reason sympathy and other inclinations are not sufficient is that they may not track the moral status of the action. A desire for someone's well-being may lead one to perform the right action in one case (as in the shopkeeper's case) but it may equally lead one to perform a morally wrong action – our desire for her well-being may move us to shelter a dangerous criminal. Thus, insofar as one's inclinations lead one to perform the right action, this is merely a lucky coincidence. Even sympathy, while it may lead one to conform with duty on many occasions, cannot do so reliably.

[sympathy] is on the same footing with other inclinations, for example, the inclination to honor, which, insofar as it fortunately lights upon what is in fact in the common interest and in conformity with duty and hence honorable, deserves praise and encouragement but not esteem; for the maxim lacks moral content, namely that of doing such actions not from inclinations but from duty. (4:398)

In contrast to both the self-interested shopkeeper and the sympathetic soul, the actions of someone who has no sympathy for her fellow man, yet does the right thing because it's morally required, do have moral worth. This suggests that the *right reason* to perform an action is that it's required by the moral law. And *to act for* the right reason is to perform it because one judges it to be required by the moral law and one has concern for morality. What is concern for morality? We can make sense of it as follows: An agent has moral concern in virtue of being disposed to perform actions that she judges to be morally required and to refrain from performing actions that she judges to be morally prohibited. There is some disagreement in the literature as to

whether the relevant disposition should be identified with desires or commitments.³⁰ Barbara Herman, for example, argues:

when we say that an action has moral worth, we mean to indicate (at the very least) that the agent acted dutifully from an interest in the rightness of his action: an interest that therefore makes its being a right action the nonaccidental effect of the agent's concern.³¹

This in turn requires that the agent “understands his action to be what morality requires” and that “this conception of his action” determines what he does.³²

For an action to have moral worth the agent must therefore judge that the action is morally required. She must then be motivated by this judgment and her (noninstrumental) interest – or desire – to do the right thing. Call this view Concern for Morality (CfM).

2. A Counterexample to CfM: Huck Finn

One problem with CfM is that it seems subject to counterexamples. Consider the following case:

Huck Finn: Huck Finn and the fugitive slave Jim are on a mission: Huck, having faked his own death, is trying to get away from his abusive father and Jim is trying to escape from slavery. Huck, a good boy, with his heart at the right spot, has never questioned the racist values of his society and so it does not occur to him that there is something morally wrong with slavery. And mid-way through his journey he comes to realize that he is about to be complicit in helping a slave escape - he is “stealing” Jim from his “owner” Miss Watson. Huck believes that this is wrong and he is haunted by conscience. He resolves to “do the right thing” and turn Jim in when the next opportunity presents himself. But sharing his adventure with Jim, his perception of Jim has shifted and he has come to see Jim as a friend and a human being like himself rather than as a slave. And when the next opportunity to turn Jim in does present himself, he finds himself unable to act on his resolution. But he doesn't revise his moral beliefs - he comes to resign himself to being a “bad boy” who doesn't have the “spunk of a rabbit”.

Huck's action seems morally praiseworthy. Huck doesn't just fail to believe that protecting Jim is the right thing to do – he falsely believes that it's the wrong thing to do. Huck is acting from,

³⁰ See Baron [1984] for arguments that it's more than a desire. See also Stratton-Lake [2003].

³¹ Herman [1981], p. 366. We find a similar thought in Stratton-Lake [2003]: “For to act from duty...is to regard myself as having sufficient reason to do some act solely in so far as I judge that it is morally required.” (p. 62)

³² *ibid.* p. 375.

what Arpaly calls, inverse akrasia. While an akratic agent acts against what is in fact her better judgment - she pours herself another Scotch when she knows she shouldn't - Huck acts against his "better" judgment. He gives in to his desire to do what's right going against his false moral belief. Thus, Huck is not saving Jim because he believes that *it's the right thing to do* - he doesn't. Nevertheless, Huck seems to deserve moral credit for his action. Hence, being motivated by a desire to do what's right cannot be necessary for moral worth. So, CfM cannot be right. Huck's case motivates an alternative account of moral worth that I will discuss in the next section.

3. Concern for Right-Making Reasons

Given that Huck does not act from concern for morality, what is it that makes his action morally praiseworthy? According to Arpaly, Huck is morally praiseworthy because he is motivated to act by those features of the situation that are morally relevant, even though he doesn't believe that they are. Arpaly argues that in order to act for the right reasons, an agent does not need to be motivated by a desire to do the right thing and she doesn't need to have true beliefs about what morality requires:

Usually, "acting from duty" and "acting for the sake of the fine" are taken to indicate not simply acting for moral reasons, but rather acting for reasons believed or known to be moral reasons — acting out of a desire to do that which is right. But moral worth is fundamentally about acting for moral reasons, not about acting for reasons believed or known to be such, and distinguishing the two is important in evaluating moral agents.³³

According to Arpaly, Huck is moved by his desire for Jim's freedom and well-being; he doesn't want Jim to suffer. And the fact that turning Jim in would lead him to suffer is what explains why protecting him is the right thing to do. In other words, Huck isn't motivated by the fact that protecting Jim *is* the right thing to do; he doesn't realize that it is. But he is motivated by the facts that *make* it the right thing to do: the fact that it saves Jim from abuse and suffering.

Arpaly suggests an alternative way of filling in the slogan "morally worthy action is action performed for the right reasons". What matters for moral worth is that the agent be motivated by concern for right-making reasons. What are the right-*making* reasons? A right-making reason is a feature of the situation that determines the moral status of the action; it's the nonmoral feature of the situation on which its moral status supervenes and to which we can appeal in order to

³³ Arpaly [2003], p. 73.

explain why the action is objectively right.³⁴ It's those features that are "*in fact* morally relevant", not those that "*the agent takes* to be morally relevant".³⁵ For example, what makes it right for Huck to protect Jim is that Jim would suffer greatly if Huck betrayed him. These factors *explain* why it would be morally wrong for Huck to turn him in. What is it for an agent *to act for* those reasons? According to Arpaly, it's to act from a non-instrumental desire for those features of the situation that are the right-making reasons. The agent need not believe that they are right-making reasons.

To say that a person acts out of moral concern is to say that a person acts out of an intrinsic (noninstrumental) desire to follow (that which in fact is) morality, or a noninstrumental desire to take the course of action that has those features that make actions morally right...³⁶

Call this account Concern for Right-Making Reasons (CfRMR).

4. Some Counterexamples to CfRMR: Moral Testimony

One undeniably attractive feature of CfRMR is that it can account for cases like Huck Finn, in which an agent acts in praiseworthy ways despite having false moral beliefs. But in this section, I will argue that there are other cases of morally worthy actions that CfRMR fails to accommodate: cases in which an agent acts on moral testimony. Agents can fail to be sensitive to right-making reasons. By relying on moral testimony, an agent can compensate for this lack. When they succeed, this is because they're determined *to do the right thing* (whatever it is). The agent is hence motivated by concern for morality. Nevertheless, intuitively, in such cases, the agent's actions do deserve moral praise. But CfRMR cannot say that.

4.1 Ignorance

Huck is not in a position to know what morality requires because he grew up in a society in which there is misleading moral evidence: misleading evidence about the moral status of slaves, etc. Huck is blameless for his moral ignorance. Because he is morally ignorant, however, he cannot rely on his concern for morality to perform the right action. But just as Huck is not in a

³⁴ More precisely, those features that provide the *best* explanation for the moral status of the action. Such features may for example include Jim's well-being. Of course the moral status of an action may well supervene on the fundamental physical facts. But the fundamental facts of physics don't provide the *best* explanation for why it would be wrong to turn Jim in – best in the sense of being simplest, having most predictive power, etc. (Just as the fundamental physical facts do not provide the *best* explanation for why some psychological or biological fact obtains.)

³⁵ *ibid.*

³⁶ *ibid.*

position to perform the right action out of concern for morality (since he has false beliefs about what morality requires), sometimes agents are not in a position to act on concern for right-making reasons – because they do not have epistemic access to the right-making reasons or because they have reason to mistrust their sensitivity to right-making reasons. Nevertheless, their actions can have moral worth. Consider the following case:

Busy Executive: Jake is an executive for Oxfam. Because of a funding crisis, he needs to cut one of two programs that Oxfam has been supporting, either X or Y. The two projects are very different and they benefit different individuals in very different ways. Jake isn't familiar with the details and he's on a very busy schedule but he needs to make a decision. He's unhappy about the cuts and he really wants to do what's right. So, he enlists his assistant, whom he knows to be a sensible and diligent man, to look into the details and tell him what he ought to do. After much investigation, his assistant tells him - rightly - to cut program X. Jake doesn't know much about the details, but he trusts his assistant's judgment and he cuts off the funding for X.

Jake seems morally praiseworthy for his action. It's not an accident that he does the right thing in the situation, cutting off the funding for program X rather than Y. But the right action is not a result of Jake's sensitivity for the right-making reasons. Jake cannot be motivated by his desire for the features of program X that explain why it's the program he ought to cut because he does not have epistemic access to what those reasons are. Rather, Jake is motivated by concern for morality – he wants to do the right thing and because his assistant is trustworthy and reliable, Jake knows that the right thing to do is to cut funding to program X.

In the case described Jake is not acting from a concern for right-making reasons, as required by CfRMR because he does not have epistemic access to what they are. But agents can act from concern for morality, rather than concern for right-making reasons, even in situations in which they do have such access. Consider a modified version of the story above, in which Jake is very familiar with the details of both funding programs but he is uncertain which one he should cut. This is because he does not know how much weight to give to the various considerations in question. He may realize that the fact that one of the programs benefits a region which is not served by other charities is morally relevant and should figure in his decision. But he may not know just how much weight he should give it. He might also worry that because of his inexperience he could mistake some features of the situation as morally relevant when they are not. In this situation, Jake may want to rely on someone who is more experienced in making

these kinds of judgments.³⁷ When he acts on this moral testimony and, rightly, cuts off funding to project X rather than Y, it's not just an accident that he performed the right action and Jake seems praiseworthy for it. But his action was not motivated by concern for right-making reasons. Rather, Jake was motivated by his desire to do the right thing and his (true) belief based on testimony that cutting program X is the right thing to do. His concern for morality motivated Jake to both seek out moral testimony and then act on it.

But suppose that Jake's advisor doesn't just tell him that cutting program X is the right thing to do but also *why* it's the right thing to do – say, because it serves an isolated community. Wouldn't Jake's action then be motivated by his concern for right-making reasons? I don't think it would be. For note that what Jake acquires when he trusts his advisor's explanation is just another moral belief: that the fact that program Y serves an isolated community is a morally relevant consideration or that, in this particular case, it outweighs the other considerations. When Jake acts on the advice, he isn't acting on concern for right-making reasons: a desire to continue serving an isolated community. Rather, he is acting on his concern for morality along with his moral belief that the fact that a program serves an isolated community is a morally relevant consideration. In other words, what motivates Jake in these circumstances does not seem to be the first-order consideration directly but rather his belief *that* it is morally relevant.

4.2 Autism

Jake relies on moral testimony because he is insensitive to the morally relevant features of the situation *on this particular occasion*. But sometimes an agent can be *systematically* precluded from having concern for some morally relevant features altogether. When such an agent compensates for her impaired sensitivity by relying on moral testimony of others, and succeeds in doing the right thing, her actions can have moral worth. In cases like this, however, CfRMR cannot explain why the agent's action have moral worth.

Much of the discussion of morally worthy actions has revolved around toy cases. Here I want to look at an example that is very closely modeled on a real one – on experiences described by Temple Grandin. Grandin, who is on the high-functioning end of Autism Spectrum Disorder (ASD) has become an influential spokesperson for autism and has written extensively about the

³⁷ See Sliwa [2012] for a discussion how agents can come to have greater moral expertise in a given domain.

ways in which her condition affects the way she navigates social interactions. ASD is a very complex condition whose manifestations and their degree greatly vary from individual to individual. Individuals with autism aren't a homogeneous group and I do not mean to suggest otherwise. Nevertheless, Grandin's experiences give us some insight into how *some* agents with autism navigate moral decisions. And it's valuable to look at an example of moral motivation that is quite different from the one of Huck Finn, which has received so much attention in the recent literature. Here's the plan: I will first argue that it is plausible that ASD affects an agent's sensitivity to *some* right-making reasons. I will then present the case and argue that when the agent compensates for her impaired sensitivity, her actions can have moral worth.

Grandin describes how ASD affects sensitivity to right-making reasons in two ways. First, ASD makes it difficult for her to recognize and interpret emotional states of others. Grandin describes that she lacks the "social sense" that "neurotypical people" rely on to guide them through both social interactions and moral decisions. An important part of this "social sense" is a capacity for empathy, which is diminished in individuals with ASD. Empathy enables us to relate to others emotionally - to read emotions of others and to respond to them.³⁸ Grandin describes:

Emotions are hard to figure out; they're not logical. My emotional make-up is simple. Everything I feel falls into one of a few categories: happy, sad, scared or angry.³⁹

Secondly, what Grandin describes as lack of "social sense" doesn't just place individuals with ASD at the same kind of epistemic disadvantage with respect to recognizing emotional states of others as nearsightedness places agents with respect to seeing far-away objects. It's not just a matter of making it harder for agents with ASD to *see* some features of their surroundings. It also puts these agents at an epistemic disadvantage with respect to recognizing these features *as morally relevant* even when they are perceived. To return to the example of Huck Finn - Huck doesn't have to deliberate about whether Jim's humanity is morally relevant and whether it provides him with a reason to help him. Huck's desires are sensitive to these features and motivate him to act, even as his moral beliefs are screwed up by his upbringing. In contrast, Grandin describes that for her making moral decisions is like solving a puzzle:

It's almost like being Sherlock Holmes, looking at every little minute detail and its relevance. But it requires that I wade through a zillion other details to find the ones that belong to the

³⁸ See Kennett [2002].

³⁹ Baron & Grandin [2005], p. 34-35.

puzzle. It's a process of picking out the relevant details. Maybe that will help people understand why responses can take longer for a person on the spectrum.⁴⁰

Despite these difficulties, Grandin *has* achieved moral autonomy. Because her concern *de dicto* to do the right thing she successfully compensates for her diminished sensitivity to some morally relevant features. Because in many situations she cannot just immediately see what to do, she relies on rules of conduct that she infers on the basis of past experiences.

By the time I was a senior in high school I had a system for categorizing some of the social rules of life [...] I reasoned that in order to maintain a civilized society, there must be prohibitions against certain highly destructive or physically harmful behaviors, such as murder, arson, rape, lying under oath, stealing, looting and injuring other people.⁴¹

Grandin reports that learning to respond in the right way to a given situation was a matter of experience and trial and error - "putting more data on her hard-drive". The more situations she had memorized, the greater the sample on which she could draw for similarities to a present situation and from which she could extract the fine-grained rule of conduct.

Another important resource is moral testimony.⁴² Moral testimony is crucial both to discover what the moral rules:

Whenever my sister Izzy and I were being rude, Mother...would immediately interject and very firmly tell us to stop that behavior and that you don't talk that way. That's all she'd say. It was the experience itself and being able to clearly categorize it as "inappropriate" that helped me learn.

And it is crucial to learning how to apply them in a given situation:

During one of the first projects I ever worked on I made the mistake of criticizing some welding...The guy's welding wasn't the greatest but that wasn't the right way to criticize it. Harley, the plant engineer brought me right up into his office and told me, real straightforward, that what I had done was wrong and that I needed to go up to the cafeteria right away and apologize to the welder.⁴³

Thus agents with autism can become very successful at making moral decisions and achieving moral autonomy. But unlike in cases such as Huck Finn's, agents in circumstances like Grandin's

⁴⁰ Barron and Grandin [2005], p. 101.

⁴¹ *ibid*, p. 130.

⁴² See also Kennett [2002], p. 253.

⁴³ Barron and Grandin, p. 203.

rely crucially on their moral beliefs. Beliefs about what the right thing to do is in various circumstances, beliefs about which features of a situation are morally relevant, how much weight they should be given, how to balance them against other morally relevant considerations play an essential role in her moral motivation. And Grandin relies on moral concern *de dicto* - she is motivated to do the right thing - whatever that turns out to be. In fact, this concern to do what's right motivates her both to go through the effort of figuring out *what* that is and then to perform the right action.

With this background, I would like to consider a particular example of an agent who, for neurological reasons, faces similar difficulties as Grandin in recognizing the emotions and weighing their moral significance. He, too, relies on rules of conduct to determine what the right thing to do is. And sometimes he finds himself in situations in which he is uncertain how to apply them:

Honesty: Arthur knows that it's important to be honest and to generally tell the truth. But now his aunt has written him an email asking him whether he enjoyed himself at the family gathering last weekend. As a matter of fact, Arthur did not enjoy himself at all. Based on previous experiences, however, he is fairly confident that saying this would upset his aunt. He's not sure whether the fact that his aunt will be upset is relevant - after all, sometimes you have to tell the truth, even if people do get upset. But he knows that *sometimes* it is relevant and then you shouldn't tell the truth. Arthur can't tell what the right thing to do is, and after much deliberation he decides to ask his friend. His friend tells him that the right thing to do in this particular situation is to avoid telling the truth.

Suppose that Arthur does take his friend's moral testimony and writes a polite email back, concealing his resentment of the family meeting and repressing his inclination to be honest about it. It seems to me that Arthur is morally praiseworthy for his action. It's not an accident that he does the right thing - he manages to do what's right because he cares about doing what's right. But Arthur is not motivated by concern for right-making reasons - he relies on moral testimony precisely because he knows that he is not sufficiently sensitive to the morally relevant features of the situation. While Arthur's desires do not track the morally relevant features of the situation, he does very much want to do the right thing - whatever this is. It's this concern for morality that

motivates him both to seek out advice and to follow it. He resists telling the truth because (he believes) this is *the right thing to do*.⁴⁴ On CfRMR then, his action doesn't have moral worth.

Note that the worry is more serious than just that CfRMR cannot accommodate this particular case. The worry is that the focus on concern for right-making reasons excludes many of the right actions of agents like Grandin or Arthur from moral praise altogether. Their condition affects a wide range of their moral decision-making. Individuals like Grandin *are* motivated to do the right thing, and they *are* successful in acting on their motivation - through much patience and effort, they succeed in compensating for their impaired sensitivity to some right-making reasons. If CfRMR cannot accommodate actions such as Arthur's in Honesty as having moral worth, this is bad news for CfRMR.

4.3 Is Concern for Morality Moral Fetishism?

I have argued that it's not an accident that Jake and Arthur succeed in performing the right action: they succeed because they have concern for morality. They want to do the right thing even when they are not in a position to know what it is. This motivates them both to seek out moral testimony and to act on it. Recently, however, the desire to do what's right - whatever it is - has been getting bad press. It's been argued that there is something morally objectionable about an agent who acts from such a desire. Michael Smith, for example, argues that good people don't care about doing the right thing, rather they care about "honesty, the weal and woe of their children and friends, the well-being of their fellows, people getting what they deserve". Someone who is motivated to act by their concern for morality rather than these concerns is someone who *fetishizes morality*. This is morally objectionable:

...commonsense tells us that being so motivated is a fetish or moral vice, not the one and only moral virtue.⁴⁵

People certainly are capable of fetishizing many things - so it may well be possible to fetishize morality. But what exactly would it be for someone to fetishize morality? And why would it be

⁴⁴ Grandin [2005], pp. 206-207 describes similar cases: "How do you know when to be completely honest and when not? That's where the ability to construct new and more refined categories is useful. Aunt Bella's [hideous] hat is not very important in the scheme of things, so the degree of honesty of my response can be lower, if needed, and I'd only offer a comment if directly asked to do so. But suppose her hat was constructed of poisonous materials that could make her sick. I'd be frank about that, because it might injure her. Her hat moved higher up the scale of importance because of the materials it was made from."

⁴⁵ Smith [1994], p. 75.

objectionable? Consider someone whom we might describe as having a fitness fetish: this would be someone who constantly obsesses about going to the gym, who talks of nothing else but protein powders and exercise regimes, who is resentful of anyone getting in the way of her gym schedule. There is plenty to dislike about fitness fetishists: they are boring. They may judge others who are less committed to working out than they are and feel superior to them. But what about acting from a concern for fitness? A fitness fetish might sometimes lead an agent to act from a concern for fitness – as when the agent buys the newest pair of running shoes because she believes that it will make her fitter. But this seems neither distinctive of a fitness fetish nor does it seem particularly objectionable. In fact, an agent who hates lifting weights is much more likely to rely on a desire for fitness to motivate himself for a trip to the gym – “I don’t enjoy it but it will keep me fit”, he might say to himself, reluctantly getting off the couch.

We can imagine someone fetishizing morality in the same way: someone who spends all their time obsessing about donating to charities and volunteering in homeless shelters, who judges others whom she regards as less devoted to doing good than she is and who gets tacit pleasure from this sense of superiority. But note that being motivated by wanting to do the right thing seems neither distinctive of a moral fetish nor particularly objectionable.⁴⁶ Neither Arthur, Jake, nor Grandin strike us as being moral fetishists. They rely on their concern for morality because they recognize that they fail to be sensitive to the right-making reasons yet they want to do the right thing.

Huck’s action reflects something admirable about him: it reflects his sensitivity to right-making reasons. But Jake and Arthur’s actions also reflect something admirable: they reflect Jake and Arthur’s sensitivity to the fact that they are morally fallible and their willingness to compensate for their limitations. A morally good agent is sensitive to both: to right-making reasons, as well as to her own fallibility in responding to them. Both are important aspects of a good moral character.

5. Moral Worth and Moral Imperfection

An ideal moral agent should be both sensitive to the right-making reasons as well as having true moral beliefs. We, however, are not morally perfect and we can fall short of the ideal in two

⁴⁶ See also Lillehammer [1997] for a defense of moral concern *de dicto* along similar lines.

different respects: we can fail to be sensitive to right-making reasons and we can have false moral beliefs. In the case of either shortcoming it's possible to compensate. Jake and Arthur compensate for their lack of sensitivity to right-making reasons by relying on moral testimony and their concern for doing the right thing. Huck's sensitivity to right-making reasons allows him to compensate for his false moral beliefs. In both cases, it's not an accident that the agents in question succeeded in performing the right action.

While the two accounts of moral worth we have seen are each compatible with one kind of imperfection, neither is compatible with both. On CfM an action can have moral worth even if the agent fails to be sensitive to right-making reasons. Hence, CfM can accommodate the cases of moral testimony that we have discussed because in those cases the agents are motivated by concern for morality and have true beliefs about what morality requires. False moral beliefs, however, are not compatible with moral worth for CfM. CfM cannot accommodate Huck Finn's action. CfRMR, on the other hand, makes moral worth compatible with false moral beliefs but not with a lack of sensitivity to right-making reasons. Thus CfRMR can accommodate Huck Finn but not Jake and Arthur.

But is an agent who has false moral beliefs thereby morally imperfect? Why should we take false moral beliefs to be a *moral failing*? Some have disputed that it is. Driver, for example, objects:

The Huckleberry Finn case is an extreme, since it illustrates a person who has demonstrably false beliefs about the good. But if virtue theorists insist that Huckleberry lacks virtue because of his flawed cognitive state, then this is bad news for most of us, who, even in some small way, are likely to harbor false views of value. The psychological requirements placed on virtue in the classical tradition seem far too rigid and unrealistic.⁴⁷ (Driver, xvi)

There are two reasons to think that Huck is morally imperfect in virtue of having false moral beliefs.⁴⁸ For one, there are some right actions that, because of his false moral beliefs, he is not in a position to perform. Given his concern for right-making reasons, Huck is in a position to do the right thing when he has to decide quickly, when there is no time to deliberate. But given his false moral beliefs, Huck is not in a good position to make decisions when he has to rely on reflection. Secondly, given his akratic state, the right action he does perform - saving Jim - is less

⁴⁷ Driver [2000], p. xvi.

⁴⁸ See Harman [2011] for an argument that moral ignorance is not just a moral imperfection but always blameworthy.

counterfactually stable than it would be if, in addition to having concern for right-making reasons, Huck also had true moral beliefs. It's for example unclear whether Huck would have protected Jim if he had known him a bit shorter and they hadn't quite made it through so many adventures together. It's also unclear if Huck would have protected Jim if he had been confronted by an adult that he is familiar with – he was able to lie to some strangers. He may not have been able to do so to his pastor, for example.

Compare Huck's case to that of someone who has true moral beliefs but whose desires are sensitive to morally irrelevant factors. Suppose that Leah, a hiring manager for a large firm, grew up in a society that is structured by race, in which income and social status are correlated with one's race. Leah does not have racist beliefs – in fact, she believes that racism is wrong and this belief guides many of her actions. She does not support politicians who take certain views on immigration reform, she votes against new voting laws that would require voters to have a photo ID. But looking at much of her behavior in other contexts, however, a very different picture emerges. In making hiring decisions for a big company, she consistently privileges members of her own race. It's not that she's doing so consciously. Rather when she looks at two applications, it just happens that the strengths of the applicant of her own race stand out more to her, while it's the weaknesses of the other application that draw her attention. As a consequence, Leah's hiring record shows a significant slant towards candidates of her own race.⁴⁹

Note that Leah is no more blameworthy for her lack of responsiveness to morally relevant features than Huck is to blame for his false moral beliefs – both are the result of growing up in an unjust society. Unlike Huck, Leah, when acting unreflectively, is disposed to discriminate against others on the basis of their race. Nevertheless, Leah and Huck's moral shortcomings are parallel in an important way: in the case of both agents moral beliefs conflict with their desires. They just differ in which one of these is on the wrong track. Just as in Huck's case this conflict implies that there are situations in which he is not in a position to do what's morally right, so there are situations in which Leah is not in a position to do what is morally right. When Leah does need to make decisions that rely on moral reflection, she is in a good position to make the right choice. We can be confident that Leah won't donate money to organizations supporting racist causes,

⁴⁹ A very similar case is described in Schwitzgebel [2010], though Schwitzgebel's main concern is whether the agent's explicitly endorsed mental state really amounts to a full belief. It seems to me plausible that in Leah's case we should ascribe her the anti-racist beliefs but I hope that nothing that I say about moral worth hinges on this.

that she will be careful not to support politicians who try to exploit prejudices for their own political advantage, and we can count on her to eloquently and convincingly argue for and defend her position when needed. To be in a position to perform these actions, it's essential that an agent has true moral beliefs and that she can reason about what the right thing to do is.

Morally imperfect agents can perform morally worthy actions – moral worth does not require virtue. Agents can fall short of moral perfection in more than one way. But importantly, whether they fall short by having false moral beliefs or by failing to be sensitive to right-making reasons, they can compensate for their limitations. When they do and they succeed in doing the right thing as a result, they deserve moral credit for their actions. The upshot is a desideratum for an account of moral worth: moral worth must be compatible with the various ways in which we fail to be morally perfect agents. This is a tricky desideratum to meet if we take moral worth to be about acting for the right reasons. The challenge is to give an account of reasons such that Arthur, who relies on moral testimony, and Huck Finn, who acts against his false moral beliefs, can count as acting for the right reasons. We will now look at an account of moral worth that attempts to do just that.

6. Reasons as Evidence

In this section I consider one more account of moral worth that has defended in recent literature by Markovits. This account aims at accommodating both Huck Finn-style cases *and* actions on moral testimony, such as Arthur's and Jake's. Markovits' account is similar in spirit to Arpaly's – she agrees that what makes an action morally worthy is the fact that the agent was motivated by right-making reasons. But it differs from Arpaly's in a crucial respect: on what the right-making reasons are. Recall that on Arpaly's view right-making reasons are the features of the situation which *explain* the moral status of the action – those features on which the moral status of the action supervenes. According to Markovits, however, right-making reasons are subjective. It follows from her account that actions can be morally praiseworthy even when the agent fails to perform the objectively right action:

The moral reasons for us to perform some action are subjective — we are morally required to do only what we have sufficient epistemic reason to believe it would be best to do, not what it would (in fact) be best to do.⁵⁰

⁵⁰ Markovits [2007], p. 219.

On Markovits' account, a feature of a situation is a right-making reason when it is *evidence* for the agent that this is the right action to perform - i.e. when it gives the agent *good reason to believe* that she should perform this action.⁵¹ Since right-making reasons are cashed out in terms of evidence, they are subjective. Evidence can, after all, sometimes be misleading. When it is misleading an agent may be acting on right-making reasons although she is doing what is objectively the wrong thing to do. Thus, in contrast to Arpaly, Herman, and Kant, on Markovits' view an action can have moral worth even when it's objectively wrong – provided the agent acted on evidence that it is the right thing to do.⁵²

Given that right-making reasons are features of the situation that are evidence for the rightness of the action, the account can accommodate cases of moral testimony. After all, moral testimony – insofar as it's reliable – just is evidence for the rightness of an action. Thus, according to Markovits, moral testimony can be a right-making reason:

If expert testimony gives us most reason to believe that some act would be best, then this testimony *is the reason* why we should perform that act [...] our agent's acting so may be *made* right by the fact that the agent advises him to do so.⁵³

Thus, as long as Arthur's and Jake's advisors are reliable, Arthur and Jake are morally praiseworthy for their right actions because they are responding to evidence that their action is right.

We have seen what right-making reasons are on Markovits' account. But what is it to act *for* those reasons? Understanding right-making reasons in terms of evidence for the rightness of the action suggests that acting on right-making evidence involves moral belief. This is typically the case when an agent acts on moral testimony: the agent forms a belief that the recommended action is the one that is morally required. But if acting for right-making reasons requires the agent to have a justified moral belief about what is morally required, then this would seem to exclude Huck Finn's action from having moral worth. After all, Huck Finn believes that in

⁵¹ Note that this differs from Arpaly's account of right-making reasons. On Arpaly's account right-making reasons are objective: right-making reasons are those features of the situation that *in fact* explain the moral status of the action. Thus, something can be a right-making feature whether or not the agent has good reason to believe that it is.

⁵² For a defense of the general claim that what it is for an agent to have a reason to φ , is for her to have evidence that she ought to φ , see Kearns & Star [2009].

⁵³ Markovits [2012], p. 24.

helping Jim he is doing the wrong thing. According to Markovits, however, for an agent to act for right-making reasons – to be motivated by evidence that an action is right – it's not necessary that she believes what her moral evidence supports, nor even that it is evidence for the rightness of one's action. When Huck is motivated by the fact that Jim is a human being very much like Huck himself, he need neither believe that this is the fact that motivates him, nor that this fact is evidence for it being right to help Jim, nor that the right thing is to help Jim. In the next section, I will raise some problems for this account.

7. Problems with Reasons as Evidence

There are two problems with Markovits' proposal. First, I argue that this account of moral worth is not extensionally adequate: while it is in part motivated by the Huck Finn case, it fails to accommodate Huck's action as having moral worth. Secondly, I argue that it counts actions as deserving moral praise that are not morally praiseworthy. The culprit, it will emerge, is Markovits' account of right-making reasons – evidence for the rightness of an action cannot be a right-making reason.

7.1 Huck Finn

On Markovits' view, Huck's action has moral worth because in helping Jim, Huck is motivated by right-making reasons. Since right-making reasons are evidence for an action's rightness, Huck's action is morally praiseworthy if and only if his evidence supports that helping Jim is the right thing to do and he is motivated by this evidence. But according to Markovits, Huck needn't actually hold the belief that is supported by his evidence. It just needs to be the case that *if* he were to believe that helping Jim is the right thing to do on the basis of his evidence, his belief would be *epistemically justified*. This means that Huck's false belief must not merely be false but *epistemically irrational*. After all, Huck's evidence supports that helping Jim is the right thing to do. Huck's belief fails to be responsive to his evidence. And if one fails to believe what one's evidence supports, one is epistemically irrational. This, however, seems implausible. In a deeply unjust society, like the one Huck grew up in, there is a lot of misleading evidence about the moral status of one's actions.⁵⁴ Of course, what exactly counts as moral evidence is a controversial question. But if moral testimony counts as moral evidence, Huck does have strong evidence in

⁵⁴ See Haslanger [2007] for a defense of the claim that part of what makes unjust social structures so insidious is that they affect what's rational for an agent to believe.

favor of turning Jim in. After all, Huck is just a boy and he has the testimony of many adults that slaves are someone's "property" and that it's wrong to "steal" them. As a child, he is justified in taking their word for many other moral propositions. And in general, they are reliable on most moral questions – other than the permissibility of slavery.⁵⁵ It is of course obvious *to us* what Huck should do. But we have the benefit of having grown up with 150 years hindsight on the question of slavery. Our epistemic situation when we consider the case is very different from Huck's: we have much more and different evidence. For Huck things are much murkier.⁵⁶ The problem with taking right-making reasons to be evidence is that Huck is only morally praiseworthy for his action if his evidence does in fact support that helping Jim is the right thing to do. And this seems, at best, questionable. This puts pressure on Markovits' proposal that acting on right-making reasons, understood in this way, is necessary for morally worthy actions.

7.2 Morally Irrelevant Factors

If right-making reasons are evidence and acting on right-making reasons is necessary for moral worth, this seems to rule out Huck's action. The account is too restrictive. But it is also too permissive: it gives us the wrong verdict about actions that clearly are not morally praiseworthy. Suppose that a white supremacist, WS, jumps into the lake to save a drowning woman. What WS is responding to is not the fact that a person is drowning but rather that a *white* person is drowning. If the woman weren't white, the WS would gleefully watch her drown instead. Of course, the fact that a *white person* is drowning is excellent evidence that a *person* is drowning. So it's excellent evidence that WS ought to jump in and save her. Hence, if what it is to be a right-making reason just is to be good evidence that something is morally required, the fact that a white person is drowning is a right-making reason. And if WS jumps into the lake because of her non-instrumental desire to save a white person, then she is acting from non-instrumental

⁵⁵ See chapter 1 for a defense of testimony as a source of moral knowledge.

⁵⁶ Could one argue that even if Huck's total evidence does not support the belief that he ought to help Jim, his action still merits praise because he is motivated by part of his evidence – the part of the evidence that is not misleading? This route is not promising. For if it's sufficient for moral worth that Huck acts on part of his evidence, it seems that he would be equally praiseworthy if he acted on the part that supports his action being wrong. Thus it seems that he would be equally praiseworthy if, swayed by the misleading testimony of the adults around him and the social practices he grew up with, he had turned Jim in. For while *we* know that this evidence is misleading, Huck doesn't; and misleading evidence is still evidence.

concern for a right-making reason.⁵⁷ So her action is morally praiseworthy. But this seems wrong. The feature of the situation that WS is responding to – being white – is morally irrelevant. Of course this feature is reliably correlated with a feature that is morally relevant – namely, being a person. But this just shows that responding to evidence that an action is right can't be enough for moral worth. You don't get moral credit when you are responding to the presence of a morally irrelevant property. But a morally irrelevant property can be evidence for the rightness of an action because it can be reliably correlated with a property that is morally relevant.

8. Conclusion

A popular approach to moral worth starts from the thought that actions are morally praiseworthy when they are *done for the right reasons*. I have considered three influential accounts of moral worth that try to develop the thought into a full-fledged account and I have argued that they fail to be satisfactory. In particular, all three fail to accommodate some cases of intuitively praiseworthy actions – they are subject to counterexamples. The discussion has also clarified some desiderata that a plausible account of moral worth should meet: moral worth must be compatible with the variety of ways that we can fall short of moral perfection. This means, it should be able to accommodate cases in which an agent non-accidentally does the right thing despite having false moral beliefs, such as Huck's case. But it should also accommodate cases such as Arthur's: cases in which agents non-accidentally perform the right action because they successfully compensate for their lack of sensitivity to right-making reasons. This motivates taking a fresh look at what makes actions morally praiseworthy – which I do in the next chapter.

⁵⁷ More specifically, Markovits takes an action to have moral worth if and only if the agent is non-instrumentally motivated by non-instrumentally right-making reasons. This doesn't help with the counterexample: the fact that a white woman is drowning is a non-instrumental reason. And WS is non-instrumentally motivated by it; she cares non-instrumentally about white people.

Chapter 3

Moral Worth and Moral Knowledge

Right action is better than knowledge, but in order to do what is right we must know what is right.

–Charlemagne–

I give money to charity to impress my mother-in-law. You give money to charity because you are troubled by the fact that we have so much and others have so little. Giving to charity is the right thing to do. Yet the two actions do not strike us as morally equal. You seem to deserve moral praise for your action, while I don't. Assessments like this are an important part of our moral practice: we don't just care about whether an agent performed the right action.

What distinguishes those actions that merit moral praise from those that do not? This is precisely what an account of moral worth tries to answer: to say that an action has moral worth just is to say that the agent deserves moral praise for it. An account of moral worth tries to identify under which conditions an agent's right actions are morally praiseworthy. We care about moral worth because we care about agents being good. In identifying what makes for morally praiseworthy actions, an account of moral worth points us to features of an agent that make for a morally good person.

Here is one observation about what seems amiss in actions that fail to be morally praiseworthy: Although they may be right, there's something worryingly flukey about them. I did give to charity to impress my mother-in-law. But had my mother-in-law been a devoted racist, I might have given money to the KKK instead. It's a fortuitous coincidence that in doing what will impress my mother-in-law I also do what's morally required. We owe the insight that such lucky coincidences are incompatible with moral worth to Kant. His conclusion is that to have moral worth actions must be performed from the motive of duty:

For, in the case of what is to be morally good it is not enough that it conform with the moral law but it must also be done *for the sake of the law*; without this, that conformity is only very contingent and precarious. (4:390)

Much of the recent literature on moral worth interprets Kant's requirement of "acting from the motive of duty" or "acting for the sake of the law" as a requirement to act "for the right reasons".⁵⁸ We have seen in the last chapter that these attempts face serious difficulties. The aim of this chapter is to develop and defend a novel account of moral worth. To do so, I take a different approach. I argue that morally worthy actions require the agent to act from *moral knowledge*:

MORAL KNOWLEDGE: An action has moral worth if and only if the agent acted from moral knowledge.⁵⁹

In the following I will spell out what moral knowledge is and what it is for an agent to act *from* moral knowledge. I show that the account is attractive for a number of reasons: moral knowledge helps us to shed light on what kind of counterfactual robustness is required for morally worthy action and explains why such counterfactual robustness obtains. It's extensionally adequate: it fits our central intuitive judgments about which actions have moral worth. It accounts for the importance of moral deliberation. And, it provides us with an attractive account of the relationship between moral worth and virtue.

The paper proceeds as follows: In *Section One*, I argue that moral knowledge can be sufficient for moral worth – provided the agent is also motivated by a desire to do the right thing. I draw out the implications of this for actions on moral testimony. In *Section Two*, I consider an apparent counterexample to my account: the case of Huck Finn. Huck Finn seems morally praiseworthy for his action despite having false moral beliefs. *Section Three* argues that moral knowledge comes in different kinds. I then show in *Section Four* that Huck Finn doesn't lack moral knowledge; he lacks a *certain kind* of moral knowledge. In *Section Five*, I show how my account provides us with an attractive explanation of why some cases of non-reflective actions are morally praiseworthy. Finally, in *Section Six*, I argue that my account avoids some difficulties that plague a recent alternative defended by Arpaly.

1. Acting on Moral Knowledge

⁵⁸ For examples, see Stratton-Lake [2003], Markovits [2007], and Korsgaard [2008].

⁵⁹ Although I am framing my proposal in terms of moral knowledge, I am also remaining largely neutral on questions about the nature of moral facts. While it may not be compatible with a moral error theory or simple boo-hooray expressivism, it is compatible with more sophisticated versions of expressivism, such as Gibbard's norm-expressivism or Blackburn's quasi-realism. Different meta-ethical positions will give different answers to the question how we can have moral knowledge. I do not hope to settle these issues in this dissertation.

Sometimes we do the right thing because we have done some thinking and we have come to the conclusion that this is the right thing to do. Consider the following case:

Whistle-Blowing: John learns that his boss is involved in some very shady transactions. He has many incentives to keep quiet about it: he has been working incredibly hard in the last couple of months to secure a promotion and he is pretty sure that whistle-blowing would cost him at least that. Even worse, it could cost him his career in the company and he really could not afford losing his job. He wants to just keep quiet about it and move on. But he finds himself increasingly restless – finally, after much agonizing, he comes to the conclusion that, much as he dreads the repercussions, the right thing to do is to speak up and report the incident. He takes a deep breath and calls the police. When later asked why he acted the way he did, he responds that “I could see that it was the right thing to do”.

John’s action *seems* morally praiseworthy. He jeopardizes his career in order to do the right thing. What is it about John’s motivation in virtue of which his action has moral worth?

It’s tempting to point to John’s desire to do what’s right. According to Herman actions have moral worth when they result from the agent’s taking an “interest in morality”. Similarly, Baron has argued that an agent deserves moral praise for an action when her action reflects a “commitment to acting from duty”.⁶⁰ It’s true that the fact that John cares about doing the right thing and that he does so non-instrumentally is important.⁶¹ If John only decided to report the incident because he wanted to impress his girlfriend, we wouldn’t give him moral credit for it. But John’s desire to do the right thing can only be *part* of the explanation for why John is morally praiseworthy for blowing the whistle on his boss.

To see that it can only be part of the explanation, consider a variant of the case. Suppose that when faced with the question whether the right thing to do is to keep quiet or to file the report, rather than deliberating, John resolves it by flipping a coin. “If the coin lands heads,” he thinks to himself, “this means that the right thing to do is to file the report. If it lands tails, it means that it’s permissible for me to keep quiet about it.” The coin lands heads and John comes to believe that morality requires him to file the report. And his belief is true. He overcomes his apprehension and does the right thing. While we may have some admiration for John’s desire to do the right thing, his action does not seem to be morally praiseworthy. John succeeds in doing

⁶⁰ See Herman [1981] and Baron [1984].

⁶¹ Isn’t concern to do what morality requires something like moral fetishism? No. I argue against this in chapter 2. See also Lillehammer [1997].

what's right. Given his strange epistemic practices, however, this success is extremely precarious. John could have very easily done the wrong thing. If the coin had landed tails, John would have formed a false belief about what is morally required of him. The problem is that a coin flip simply isn't a reliable guide to what morality requires. If he is lucky, John may end up with a true belief about what the right thing to do is. Even if he does, however, John doesn't *know* what the right thing to do is. Morally worthy action requires moral knowledge.⁶²

One reason for taking moral knowledge to be necessary for moral worth is that it can explain the counterfactual robustness of morally worthy actions. In Whistle-Blowing, John's right action seems praiseworthy because it doesn't seem that he did the right thing only accidentally. He would not deserve moral credit for reporting the incident, if he would have reported it even in circumstances in which it was clearly the wrong thing to do – when doing so, for example, would have led to the death of many innocent people. The counterfactual robustness in question is not unique to moral praise. In most nonmoral domains we give praise for success only when the success is not a mere fluke. Consider the doctor who cures my headache by prescribing me some aspirin. He doesn't deserve medical credit for the cure if aspirin is the only medication he ever prescribes and he prescribes it to all patients, no matter what ailment they come in for – he falsely believes that aspirin is the cure for everything and anything. Nor does he deserve credit if he decides which medicine to prescribe by throwing a die. In either case he doesn't deserve medical credit even if his treatment is successful and even if he genuinely desires to only prescribe what will help the patient. In fact, he doesn't deserve credit for his action even if due to strange circumstances, he is justified in believing that aspirin is a cure for all conditions or that throwing a dice is a reliable decision-procedure for prescriptions. Good intentions to help the patient aren't on their own sufficient to bring about success; nor are they sufficient together with justified true belief about how to help the patient. What's missing is medical knowledge – the doctor must *know* which medicine to prescribe the patient. The same goes for the moral case: the desire to do what's right is not enough for moral worth even when conjoined with a justified true belief about what morality requires. The agent needs to have moral knowledge.

⁶² This case illustrates that justified belief is not enough either. We can imagine a scenario in which John may very well be justified in thinking that the requirements of morality reveal themselves through coin tosses and in which he therefore is justified in forming his belief on the basis of the coin toss. Still, John is lucky that he managed to do the right thing.

On my proposal, if John's action in Whistle-Blowing is morally praiseworthy, John must have moral knowledge. He must know that reporting the incident is the right thing to do, not just truly believe that it is the right thing to do, even if this belief is justified. How did John acquire this knowledge? As I have described the case, John starts out uncertain about what the right thing to do is; he makes up his mind by deliberating. Thus the claim that moral knowledge explains why John's action has moral worth is only plausible if it's plausible that deliberation can be a source of moral knowledge.

1.1 Deliberation as a Source of Moral Knowledge

The proposal that moral deliberation can be a source of moral knowledge is very much Kantian in spirit. Kant himself held that it is *necessary* to engage in moral reflection to know what morality requires of us. We have an ability to reason specifically in order to *be good*:

the true vocation of reason must be to produce a will that is good, not perhaps as a means to other purposes, but good in itself, for which reason was absolutely necessary. (4:396)

But, according to Kant, of course not any kind of deliberation will do – to get at what morality requires of us we have to deliberate in the right way. In particular, we must be careful that we don't let ourselves be biased by our inclinations. On Kant's view, our desires are no guide to what morality requires of us. This is because our desires aim not at right action but at our survival and happiness. For this reason, when deliberating about what the right thing to do is, we ought to pay as little attention to what we are inclined to do as possible, lest our desires lead us to rationalize actions to which we are drawn by our self-interest.⁶³

This is precisely why the Categorical Imperative is so valuable. It provides us, Kant argues, with an idiot-proof method for coming to know what morality requires of us, even when we are faced with novel situations:

I do not...need any penetrating acuteness to see what I have to do in order that my volition be morally good. Inexperienced in the course of the world, incapable of being prepared for whatever might come to pass in it, I ask myself only: can you also will that your maxim become a universal law? If not, then it is to be repudiated... (4:403)

⁶³ As Kant argues: "The human being feels within himself a powerful counterweight to all the commands of duty [...] – the counterweight of his needs and inclinations, the entire satisfaction of which he sums up under the name happiness. [...] But from this there arises a natural dialectic, that is, a propensity to rationalize against those strict laws of duty and to cast doubt upon their validity, or at least upon their purity and strictness, and, where possible, to make them better suited to our wishes and inclinations." (4:405)

The Categorical Imperative allows us to come to know what we ought to do in situations in which we we might otherwise risk giving too much weight to our desires. Many contemporary Kantians – Barbara Herman and Christine Korsgaard, for example – agree with this picture. They take moral deliberation to be crucial for moral worth because unless we pause and reflect on what morality requires of us, our default is just to act on our desires. Those, however, are not reliable guides to what we morally ought to do.

But moral deliberation does not just seem valuable as a corrective to potentially self-interested desires. Just as in nonmoral domains, moral deliberation can be a useful tool for resolving uncertainty about which course of action to take. Consider a doctor who is uncertain which one of two treatments to prescribe a patient. She may be able to resolve her uncertainty by consciously reflecting on similar cases she has seen in the past, she may reflect on what she has learned in medical school, she may weigh the possibilities of side-effects against the success rates. For medical questions, deliberation can take very many forms. Similarly, it seems plausible that moral deliberation can take many different forms. It needn't always involve evaluating whether the course of action is universalizable. In Whistle-Blowing, John may reflect on his past experiences and compare his present situation to ones he has encountered in the past, he can try to imagine what people he admires and respects would do if they were in his situation, he can reflect on the moral principles he endorses, he can think about how he would justify taking one course of action to his children, he can try to simulate what it would be like to decide one way or another... There are many ways of deliberating about what the right thing to do is and it seems plausible that at least some of those can be reliable – just as there are many reliable ways to deliberate about nonmoral questions. Insofar as it is reliable, it seems plausible that deliberation can be a source of knowledge – in moral just as in nonmoral domains.

1.2 Moral Worth and Moral Testimony

Moral worth, I have argued, requires moral knowledge – merely true belief, even if it's justified, does not suffice. And I have argued the proposal can accommodate our intuitive reaction in Whistle-Blowing. Moral deliberation often can be a source of moral knowledge. In Whistle-Blowing it seems plausible that by engaging in moral reflection, John has come to know that the right thing to do is to file the report. When he is motivated both by knowing what the right

thing to do is and his desire to do the right thing, he acts from moral knowledge. But can moral knowledge (along with a desire to do the right thing) be sufficient for moral worth? In this section I argue that it is.

If what matters for moral worth is moral knowledge and, as I argued in the last section, moral deliberation often is a source of moral knowledge, then moral deliberation has an important role to play in morally worthy action. But, deliberation may not be necessary for moral worth. Contrary to Kant, there may be other ways in which an agent can acquire moral knowledge. Suppose, for example, that in Whistle-Blowing, John chooses to resolve his uncertainty about what the right thing to do is in a different way. Imagine that instead of deliberating, John seeks out the moral advice of someone whom he knows to have good moral judgment and much more experience with these kinds of situations than John himself. Suppose that this advisor tells John, rightly, that the right thing to do for John is to report the incident. John trusts the advisor, acts on this advice and reports the incident. In this case, John relies on moral testimony to determine what he ought to do; he doesn't deliberate. Still, John's action does seem morally praiseworthy. The intuition that John's action is morally praiseworthy in this case is in part based on our judgment that it's not just a fluke that John succeeded in doing what morality required. And there's good reason to think that this judgment is right: insofar as John's advisor is trustworthy and reliable, John can come to have moral knowledge on the basis of her testimony: John's right action isn't just a lucky accident because John *knows* what the right thing to do is.⁶⁴

In fact, there are many situations in which the agent may not be in an epistemic position to acquire moral knowledge by deliberating. In such cases moral testimony can be the *only* source of moral knowledge available to the agent. This may be because the agent is under time pressure and simply does not have the time to think through the question herself, or because she simply doesn't trust her own reasoning on the matter; she might be worried about being biased or distracted. In such cases it seems the agent *shouldn't* deliberate; she should seek out moral advice instead. If she insists on thinking through the question herself, despite, for example, having reason to believe that she is biased, her actions do not seem praiseworthy even if she does in fact manage to do the right thing. A plausible explanation of this is that in such cases the agent fails

⁶⁴ Aren't there special problems about acquiring moral knowledge on the basis of moral testimony? No. If you're not willing to take this on philosophical testimony, see chapter 1 for an argument.

to have moral knowledge; her knowledge is *defeated*. Moral deliberation then not only isn't necessary for moral worth, it can sometimes be incompatible with it.

MORAL KNOWLEDGE accommodates our intuition that John is morally praiseworthy for filing the report even when he does not deliberate but relies on moral advice instead. And it explains why his action is morally praiseworthy; it explains why it's not accidental that John did the right thing. Such cases suggest that moral knowledge along with the desire to do the right thing can be sufficient for moral worth. This, however, is controversial. Korsgaard and Hills have argued that moral knowledge cannot be sufficient for moral worth, since this would make morally praiseworthy action *too easy*; it results in an account that's too inclusive. The argument specifically targets actions based on moral testimony. They argue that there are cases in which moral testimony leads an agent to have moral knowledge. Yet the agent does not seem to deserve moral praise for her actions. They conclude that an agent needs something more than mere knowledge – moral understanding, for example – in order to deserve moral praise for her actions. Moral understanding, however, cannot be simply acquired by someone telling you what the right thing to do is.

Hills presents us with the following counterexample to the sufficiency claim of MORAL KNOWLEDGE:

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.⁶⁵

Hills argues that Ron knows that the right thing to do is to refrain from killing Tamara on the basis of his rabbi's testimony. Nevertheless, when Ron refrains from killing Tamara because he wants to do the right thing and he knows that the right thing is to let her live, Ron's action is not morally praiseworthy. Hence, moral knowledge cannot be sufficient for morally worthy actions – rather the agent needs to have moral understanding. Furthermore, moral testimony cannot be the basis for morally worthy actions, because while it may be a source of moral knowledge, it cannot lead an agent to have moral understanding.

⁶⁵ Hills [2009], p. 115. Hills' case is similar to one discussed in Arpaly [2003]. A similar point against moral testimony is also raised by McGrath [2010].

Similarly, Korsgaard argues:

to act from duty is not just to be moved by a blank conviction that an action is required, but rather to be moved by a more substantial thought which inherently involves an intelligent view of why the action is required.⁶⁶

Why require the more demanding condition? Korsgaard is worried that the “blank conviction that an action is required” is something that an “ordinary conscientious person” may well share with “the Nazi soldier who thinks of “duty” as carrying out the orders of his superiors”.

These worries, however, are misguided. On closer look, these arguments against the sufficiency of moral knowledge turn out to be unsound: the agents described may have a true belief on the basis of the testimony but they do not have moral knowledge. Consider Hills’ case. Even if the rabbi’s testimony is reliable, it doesn’t follow that Ron knows that killing Tamara would be the wrong thing to do. Note that there is something Gettier-ish about the case. It implies that after hearing the rabbi’s testimony Ron believes that killing Tamara would be wrong *because Tamara is a fellow Jew* – not because she is a fellow human being. Ron’s belief that killing Tamara would be wrong seems to be partially based on a false moral belief. And that seems enough to defeat the knowledge attribution, even if the rabbi’s testimony was reliable. And in Korsgaard’s example, she glosses over an important difference between the “blank conviction” of the ordinary conscientious person who follows the commands of her conscience and that of the Nazi soldier who follows the orders of his superior – those of the former may well qualify as moral knowledge. Those of the latter most certainly will not. Given just how wrong the Nazi soldier is about some fundamental moral truth, he is not in an epistemic position to acquire moral knowledge.

The problem with Hills and Korsgaard’s argument isn’t just that the cases they happened to come up with aren’t convincing. There’s a more general issue here: the reason they take moral knowledge to be too easy a condition to meet for moral worth is that they assume that moral knowledge is very easy to get. In particular, they assume that knowledge – both moral and nonmoral – by testimony is extremely cheap: as long as the testifier is reliable and her testimony is true, the hearer gains knowledge. But this is not enough: even when the testifier is reliable and her testimony is true, the testimony may fail to transmit knowledge if the recipient fails to meet certain epistemic conditions. Spelling out what exactly these epistemic conditions are is

⁶⁶ Korsgaard [2008], p. 182.

unfortunately beyond the scope of the paper.⁶⁷ But so much seems plausible: an agent cannot come to know that *p* on the basis of testimony if she is too credulous, i.e. if she is completely unable to distinguish a reliable from an unreliable source of testimony about *p*. Suppose that an agent would believe what you tell her about the outside temperature, even if you tell her that it's 300 degrees Fahrenheit. Such agent just doesn't seem to be in an epistemic position to come to know the outside temperature even when the testimony she does receive is in fact reliable and true. But this is exactly the epistemic position that Ron and the Nazi soldier are in – both are mistaken about some very fundamental moral truths. Even when they receive testimony that is true and reliable, their belief hence cannot meet the conditions for knowledge. This is very different from John in Whistle-Blowing who seeks out moral advice because he has reason not to trust his own judgment on the matter. John isn't radically mistaken about what morality requires; he just doesn't know what the right thing to do is *in this particular case*. He wouldn't go along with advice to execute babies or burn down synagogues.⁶⁸ Moral testimony is not valuable because it makes moral knowledge *easy*; it's valuable because sometimes it is *easier* and sometimes it's the *only way* to get moral knowledge.

I have argued that some cases in which an agent succeeds in doing the right thing by relying on moral testimony rather than by deliberating on her own can have moral worth. MORAL KNOWLEDGE allows us to accommodate this intuition. This is because moral testimony can give an agent moral knowledge and, on my proposal, moral knowledge (together with the desire to do the right thing) is sufficient for moral worth.

2. Is moral knowledge really necessary?

I have argued that moral knowledge along with concern for doing the right thing is *sufficient* for morally worthy actions and that this does not make morally worthy actions *too easy*. In this section I respond to the opposite worry: that MORAL KNOWLEDGE makes moral worth *too hard*. This is because it claims that moral knowledge is *necessary* for moral worth. This has struck many as implausible. The requirement seems at odds with our intuitions about cases like the one of Huck Finn, in which an agent seems morally praiseworthy for her action even as she holds false

⁶⁷ But for an excellent discussion of the epistemic virtues required to acquire knowledge by testimony, see Fricker [2007].

⁶⁸ See Chapter 1 for a discussion of how we can sometimes reasonably regard someone as more reliable on a particular moral question than ourselves.

moral beliefs. Huck Finn's case in particular has figured prominently in the recent literature on moral worth; it's taken as conclusive evidence that moral knowledge cannot be necessary for moral worth. Here is the case:

Huck Finn: Huck is helping the fugitive slave Jim to escape from slavery. Huck is a good boy but having grown up in the Deep South of the 19th century, he has absorbed the false moral beliefs about slavery and the moral status of black people that are prevalent in his society. Mid-way through his journey it dawns on him that what he's doing is "wrong" - he is "stealing" Jim from his "owner" Miss Watson. Haunted by his conscience, Huck resolves to "do the right thing" and turn Jim in when the next opportunity presents itself. But when the opportunity comes, he finds himself unable to act on his resolution and he makes up an elaborate lie to protect Jim. What he doesn't do, however, is revise his moral beliefs - he concludes that being moral is just "too hard" and comes to resign himself to being a "bad boy" who doesn't have the "spunk of a rabbit".⁶⁹

Huck does seem morally praiseworthy for protecting Jim in these circumstances. Yet it seems quite clear that Huck *does not know* that protecting Jim is the right thing to do. For one, the way the case is usually described, Huck doesn't *believe* that it's the right thing to do.⁷⁰ In fact, Huck falsely believes that in helping Jim he's doing the wrong thing - that he ought to turn him in instead. Plausibly, knowing that p is incompatible with falsely believing that not-p. Secondly, there's good reason to think that Huck couldn't know that saving Jim is the right thing to do even if he didn't falsely believe otherwise. Given Huck's upbringing, he simply isn't in an *epistemic position* to know. After all, Huck is just a child and he grew up in a society in which racism is deeply entrenched and structures much of every-day life. And one of the pernicious aspects of such a society is that it provides plenty of misleading moral evidence: evidence that race is morally relevant, that slavery is permissible, etc. This means that Huck has *misleading evidence* about the moral status of his action. He has, for example, testimonial evidence that a slave is just someone's property and that helping him escape is like stealing. This is testimonial evidence by people whom he trusts and whom he has good reason to regard as reliable on moral matters. And in fact, they generally are reliable on moral matters - with the exception of questions related to the permissibility of slavery. Note also that, at least as far as Huck knows, whether it's wrong to help a slave escape is not a controversial question. There is wide-spread agreement that it's wrong and Huck is aware of this. Huck simply doesn't have the epistemic resources to challenge the

⁶⁹ This case appears in Twain (1884) in Chapter XVI.

⁷⁰ See, for example, Markovits [2007], p. 208, Arpaly [2003], Driver [2000]. See Holton [ms] for a different reading of the case.

consensus. He doesn't have the concepts necessary to name and understand the moral mistake – racism – that his contemporaries are making. Neither, for that matter, does Jim; he is subject to 'hermeneutical injustice'.⁷¹ Unjust social structures are epistemically pernicious: they generate misleading evidence that justifies beliefs which make it rational to uphold the structure.⁷²

3. Practical and Theoretical Moral Knowledge

Huck is morally praiseworthy for protecting Jim even though he fails to know that protecting Jim is the right thing to do. Doesn't this establish that moral knowledge cannot be necessary for moral worth? I will argue that it doesn't. Rather, what it establishes is that a *certain kind* of moral knowledge is not necessary for moral worth. Huck lacks *theoretical* moral knowledge that helping Jim is the right thing to do: he doesn't believe that helping Jim is the right thing to do and he lacks justification for such a belief. However, Huck does have a different kind of moral knowledge: *practical* moral knowledge. And his practical moral knowledge is crucial in explaining why he deserves moral praise for his action. Before making the case that Huck Finn has practical moral knowledge, however, we will need to get clear on the distinction between *practical* and *theoretical* moral knowledge. This is the aim of this section.

3.1 What's the Difference between Practical and Theoretical Knowledge?

When discussing practical knowledge, the focus tends to be on sports – riding a bicycle, skiing. But agents have practical knowledge in many more domains. Almost any kind of expertise involves practical knowledge. Take an experienced surgeon, for example. Being an experienced surgeon requires different kinds of knowledge. A surgeon knows what the symptoms of a ruptured appendix are, she knows the names and locations of various arteries and she knows the results of various recent scientific studies on different surgical procedures. When we say that she knows these things, we mean that she has true justified beliefs about these subject matters; she has *theoretical* knowledge. When asked, she could tell you what the symptoms are or what the study established. But being a surgeon also seems to involve knowledge of a different kind. It involves, for example, knowing how to perform an appendectomy. This isn't just a matter of having memorized the relevant textbook chapter. An experienced surgeon has the *ability* to successfully perform the procedure under a variety of circumstances and complications. To have

⁷¹ See Fricker [2007].

⁷² See Haslanger [2007] and [2010].

the ability to perform an appendectomy, you have to have the ability to hold the scalpel at the right angle, to make the incision in the right place, to identify the appendix, to make the cut, to stop the bleeding,... *Right* angle, *right* place is crucial here: to say that a surgeon has practical knowledge of how to perform an appendectomy is to say that she is sensitive to the medically relevant features of the patient and that she is able to modify the way in which she performs the appendectomy in response to those features. All this is *practical* knowledge.

In the following, I am going to presuppose a particular account of how practical knowledge differs from theoretical knowledge.⁷³ Some have denied that there are two importantly different kinds of knowledge; they point to linguistic evidence in support of their position.⁷⁴ But there are a number of reasons for thinking that they do differ. Knowing how to ski just seems intuitively different from knowing that Napoleon was defeated at Waterloo. Such a distinction is also suggested by findings in cognitive science.⁷⁵ And the linguistic evidence to which opponents of the distinction point is far from unequivocal.⁷⁶ I cannot possibly do justice to the arguments on both sides here. Thus, I will simply presuppose that a good case for taking practical knowledge to be importantly different from theoretical knowledge can be made. I will also make some assumptions about the nature of the distinction. Theoretical knowledge, I assume, is closely linked to belief: an agent's having theoretical knowledge that *p* entails that she believes that *p*. Practical knowledge, on the other hand, is closely connected to ability: an agent's having practical knowledge how to perform an action *entails* that the agent has the ability to perform it. The ability in question is one that's *internal* to the agent. A surgeon who has doesn't have a scalpel still has the ability to perform appendectomies; the external conditions just aren't such that she could exercise her ability. A (former) surgeon suffering from advanced Alzheimer no longer has the ability to perform appendectomies, even if she did have access to all the necessary tools.⁷⁷ The assumption that practical knowledge entails having an ability is controversial but I cannot argue

⁷³ A view closest to the one I am assuming has been defended in Glick's [2009]. See also Hawley [2003] for a similar account. Wallis [2008] offers an account of practical knowledge in the same spirit.

⁷⁴ See, in particular Stanley and Williamson [2000], Stanley [2011a] and [2011b]

⁷⁵ Wallis [2008] provides a good overview. See also Adams [2009], and Glick [2011] for discussion of the empirical results.

⁷⁶ See Glick [2009] for a discussion of semantic considerations that point in favor of the distinction.

⁷⁷ See Hawley [2003] and Glick [2009] for a discussion of the abilities in question.

for it here.⁷⁸ I will briefly sketch the view but my main concern is its implications for moral knowledge. However, insofar as the view I'm presupposing does have plausible implications for the moral case, this may be taken as additional support for it.

On the view I am assuming, the following are typical of theoretical knowledge: theoretical knowledge involves *belief* – saying that an agent has theoretical knowledge that *p* entails that she believes that *p*. It also involves *justification*. It tends to be linguistically accessible to the agent; typically, if I know that Bob is in his office, I can tell that he is. What theoretical knowledge is available to an agent depends on the concepts possessed by the agent – I cannot know that Bob is in his office if I do not possess the concept “an office”. And typically, the content of an agent's theoretical knowledge is available for conscious reasoning and drawing inferences – I can use my theoretical knowledge that Bob is in the office to infer that he's not at home. Correspondingly, testimony and reasoning are important sources of theoretical knowledge.

Practical knowledge, on the other hand, is closely tied to *ability*. An ability to juggle or to ride a unicycle or perform an appendectomy isn't just a matter of having some particular belief. Rather, an agent has practical knowledge how to ride a unicycle when, given the appropriate set of conditions, she *can* successfully ride a unicycle. Consequently, justification seems irrelevant for practical knowledge: whether an individual has acquired justified beliefs seems irrelevant to whether she has learned how to ride the unicycle.⁷⁹ Nor is an agent typically in a position to articulate the content of her practical knowledge – we cannot typically report *how* we ride a unicycle.

I want to get one potential source of confusion out of the way: When speaking of practical knowledge, we generally say that the agent has *knowledge how* to perform an action. But not *all* know how entails ability.⁸⁰ We might, for example, say that a judge at an ice skating competition knows how to perform a triple lutz. We could say this even if she isn't actually able to execute it (she has never skated in her life) but she has justified, true beliefs (and whatever else it takes)

⁷⁸ Again, I refer the reader to Wallis [2008], Adams [2009], and Glick [2011] and [2009]. Stanley and Williamson [2000] deny that practical knowledge involves abilities.

⁷⁹ Glick [2011], pp. 408-409.

⁸⁰ Some have concluded from this that the English expression “A knows how to φ ” is ambiguous. For a defense of the ambiguity thesis, see Glick [2009].

about how to correctly execute the jump. In this case, the judge has *theoretical* knowledge how to perform a triple lutz. On the other hand, we can imagine the skater who is able to perform the triple lutz perfectly but has no insight in how she does it. It seems right to say that the skater knows how to perform a triple lutz. But her knowledge how is *practical* – she has the ability to execute the jump. And someone who is both an excellent skater and an excellent coach has both practical and theoretical knowledge how. To avoid confusion I will always be explicit about what kind of knowledge how I refer to: I will use the terms “practical knowledge how” and “theoretical knowledge how”.⁸¹

We have seen how practical and theoretical knowledge differ. But there are also important similarities. On both theoretical and practical knowledge the agent’s “getting it right” is counterfactually robust to a certain degree. In the case of theoretical knowledge the “getting it right” refers to the agent’s *true belief*. In fact, theoretical knowledge is often taken to require “safety” and “sensitivity”: an agent who knows that *p* wouldn’t have formed the belief that *p* in certain counterfactual scenarios in which *p* were false and she would have formed the belief in different counterfactual scenarios in which *p* were also true. In the case of practical knowledge the “getting it right” refers not to true belief but to *successful action*. To say that an agent has practical knowledge how to ski is to say that she can successfully make it down the hill under a range of conditions. To say that she made it down the hill because she exercised her practical knowledge how to ski is to say that she would have made it down the hill even if it were a bit steeper, or even if she would have started from the left corner rather than from the right corner of the slope. Theoretical and practical knowledge both require sensitivity and responsiveness to relevant features in the agent’s surroundings. Theoretical knowledge requires that an agent’s beliefs are responsive to her evidence. Practical knowledge requires that the agent’s performance of the action are responsive to the features that determine its success.⁸²

⁸¹ Another important clarification: The question whether there is a distinction between practical and theoretical knowledge is independent of whether practical knowledge is, just like theoretical knowledge, propositional in some sense. While the latter is hotly debated in recent literature, the former is far less controversial. In particular, even those who argue that practical knowledge is propositional, still maintain that there is an important distinction between practical and theoretical knowledge. For example, according to Stanley and Williamson [2001] both theoretical and practical knowledge are propositional but practical knowledge involves being acquainted with the proposition under a “practical mode of presentation”. For the purpose of this paper I remain neutral on whether practical knowledge is propositional or not.

⁸² See Hawley [2003] for an attempt to spell out the counterfactual robustness in question. See also Wallis [2008].

3.2 Practical Moral Knowledge

With a rough distinction between practical and theoretical knowledge in hand, let's return to the moral domain. Theoretical moral knowledge seems rather straightforward: we have seen an instance of it in Whistle-Blowing. To act from theoretical moral knowledge is to act from a desire to do the right thing along with theoretical moral knowledge of what the right thing is. But what is *practical moral knowledge*? And what is it to act *from* practical moral knowledge?

Most generally put practical moral knowledge is closely tied to the ability *to do the right thing*. Of course this doesn't mean that it's a matter of having the ability to perform some *one particular* kind of action. Practical moral knowledge isn't like practical knowledge how to juggle or ride the unicycle. It is more like practical *medical* knowledge. Practical medical knowledge involves very many different abilities: the ability to perform appendectomies, to take someone's blood pressure, to put in an epidural. To be a competent doctor an agent needn't have all of them – though she does need to have some subset of them. Similarly, practical moral knowledge involves many different kinds of abilities: the ability to resolve a dispute fairly, the ability to treat someone with the respect they deserve, the ability to do what's just... Just as in the medical case, being a competent moral agent is compatible with falling short on some of those. And just as performing an appendectomy is a complicated procedure that requires a whole host of other abilities – the ability to hold a scalpel, the ability to isolate the appendix, to stand at the right distance from the operating table – so, for example, being able to resolve a dispute fairly may also involve many other abilities: the ability to listen without prejudice to both sides, the ability to weigh up the conflicting interests, to just name a few.

What is important is that practical moral knowledge is not just the ability to make the right moral judgments – it's not just the ability to reliably form true beliefs about what the right thing to do is. Rather, practical knowledge is the ability to *do the right thing*. This is not to deny that an agent may sometimes have the ability to acquire theoretical moral knowledge *because* she has practical moral knowledge. An agent who has practical knowledge how to treat others fairly may sometimes be in a position to reliably form true beliefs about how to treat others fairly. But she may not – she may, for example, lack the requisite concepts. Again this is not specific to the moral domain. Practical knowledge how to perform an appendectomy isn't just the ability to form true beliefs about how an appendectomy should be performed. Rather, it's the ability to

perform appendectomies. This is not to deny that an agent may be able to reliably form true beliefs about how an appendectomy should be performed because she has the ability to perform them.

The distinction between the ability to do something and the ability to form true beliefs about how to do it is very important for the question of what it is to act *from* practical knowledge. We need to distinguish between these two abilities to avoid one particularly tempting misconception about what exercising one's practical knowledge involves. According to this misconception, an agent exercises her practical knowledge how to perform an appendectomy by making a judgment about and thereby coming to believe how the appendectomy should be best performed. She then *chooses* to act in light of her best judgment. This thought is particularly popular for practical *moral* knowledge. Herman, for example, argues that only when an agent acts in light of her best moral judgment can her action have moral worth:

For an action to have moral worth, moral considerations must determine how the agent conceives of his action (he understands his action to be what morality requires), and this conception of his action must then determine what he does.⁸³

We find a similar line of thought in Korsgaard:

An agent's motivation to act involves two things — the incentive that presents the action along with its end as eligible, and the principle of volition that governs the agent's choice to act on that incentive [...] moral value rests specifically in the principle of volition that is exercised in the choice of the action. Moral value supervenes on choice.⁸⁴

The thought seems to be that while practical moral knowledge can lead an agent to see what the right thing to do is, she still has to make a choice to then act *from* her practical moral knowledge.

This, however, is not the right way to think about practical knowledge. While practical knowledge can enable an agent to make true and reliable judgments about which action she should perform, the agent doesn't need to actually make the judgment in order to exercise her practical knowledge. In fact, one reason why practical knowledge is so valuable is that it can enable us to act in the absence of such judgments, thus allowing us to put our cognitive capacities to other uses. An experienced driver doesn't need to form a conscious belief that she should shift

⁸³ Herman [1981], p. 375.

⁸⁴ Korsgaard [2008], p. 181.

into a lower gear and then decide on it. She does so unthinkingly, perhaps not even noticing it, while thinking about the song on the radio or about what to make for dinner. Similarly, an experienced surgeon doesn't *need* to consciously reflect on the right place to make the incision, or how to respond to a sudden bleeding. She is able to just respond to the circumstances without a conscious belief about what she ought to do, perhaps while consciously deliberating about some other aspect of the surgery: whether she should call in another specialist, for example. The fact that the surgeon and the driver don't make a judgment about what they ought to do and that they don't decide to act on that judgment doesn't mean that they are acting from "mere habit" or "from inclination" and that it's mere luck that they got it right. The agents are sensitive and responding to the relevant features – the medical condition of the patient or the road.

Practical moral knowledge enables the agent to do the right thing without conscious reflection on what the morally relevant features are and without having to make a conscious decision to perform the action because it's the right thing to do. This gives us a satisfying explanation for why some non-reflective actions deserve moral praise – actions in which the agent appears to non-accidentally do the right thing in the absence of a moral judgment altogether. Consider, for example, some very mundane actions for which we give agents moral credit: simple everyday niceties like holding the door for someone struggling with heavy boxes, grabbing the back of the stroller as the parent is struggling up the stairs, giving up one's seat to an elderly person...⁸⁵ Often these kinds of actions are the right thing to do – *morally* and not just as a matter of etiquette. We give agents moral credit for performing them. But it doesn't seem plausible that they are praiseworthy in virtue of the agent's having theoretical moral knowledge. In many of these cases the agent performs the right action without making a moral judgment and then deciding to act on it; she may not even consciously register that she is holding the door for someone.

When we consider these actions morally praiseworthy, it is because we don't think it's just a fluke that the agent did what was morally right in this case: we think that the agent wouldn't have held the door open if it was clearly the wrong thing to do. For example, we don't think that the agent would have held open the door for someone whom she has just seen mugging an elderly lady. But this is just to say that we take the agent to be praiseworthy for her action because we attribute to

⁸⁵ Thanks to Rae Langton for drawing my attention to these kinds of cases.

her an ability to do the right thing. This seems like a reasonable assumption: the ability in question is the kind of ability that most morally competent agents have.

How is this possible? How can an agent be responsive to all these features relevant to such complicated actions like performing an appendectomy without consciously reflecting on them? This is a question for psychologists to answer. In fact, psychologists have identified a number of mechanisms that enable an agent to be responsive to features of her environment without having to make conscious decisions how to act. Giving a full account is far beyond the scope of this paper. But one important mechanism is something like stereotyping: recognizing novel situations as relevantly similar to ones encountered in the past. The agent hence needn't consciously reflect about what the correct course of action is and she need not make a conscious choice to perform those actions – she can immediately respond to the relevant features of the situation.⁸⁶

4. Does Huck Finn have Practical Knowledge?

The last section has given us a rough idea of the distinction between practical and theoretical knowledge and how it applies to moral knowledge. With this in place, let's return to the case of Huck Finn. Huck Finn looked like a problem for the view I have been defending. On my proposal moral knowledge is a necessary condition for moral worth. Yet Huck Finn seems morally praiseworthy for protecting Jim despite having false moral beliefs about the moral status of his action. So, moral knowledge cannot be necessary for moral worth. In this section I argue that the case does not undermine the necessity claim – it supports it. This is because while Huck lacks theoretical moral knowledge, he does have practical moral knowledge. And his practical moral knowledge is crucial for evaluating whether his action has moral worth: we take his action to have moral worth only insofar as it seems plausible that Huck acted from practical moral knowledge. What Huck's case shows then is that the necessary condition of moral knowledge can be met in virtue of the agent's having either practical *or* theoretical moral knowledge.

4.1 Counterfactual Robustness

It's hard to deny that Huck's action strikes us as morally praiseworthy. It strikes us this way because it seems that Huck didn't just perform the right action by accident. This intuition – that it's not just a fluke that Huck made the right call in the situation – is in fact well supported by

⁸⁶ For an overview of these mechanisms, see Bargh & Chartrand [1999]. See also Klein [1998] and Holton [2009].

the story. It's important to note that the episode on which I focus here isn't just an isolated incident of Huck's doing the right thing – treating Jim fairly– in the narrative. There are a number of incidents that precede this particular action. I'll mention just two: one towards the beginning of the story and one later on. Huck, having faked his own death, escapes on a little island where he meets Jim, who ran away from Mrs. Watson when he overheard her thinking about selling him on to New Orleans. Jim was devastated by this prospect: it meant harsh treatment as well as being separated from his family. After sharing breakfast with Jim, Huck asks Jim why he is hiding on the island. Jim is uneasy and he gives in to Huck's questions only after Huck promises that he won't tell anyone. He reveals that he has run away and implores Huck to keep the promise: "But mind, you said you wouldn't tell – you know you said you wouldn't tell, Huck." Huck responds: "Well, I did. I said I wouldn't, and I'll stick to it. Honest *injun* I will. People would call me a low-down Abolitionist and despise me for keeping mum – but that don't make no difference."⁸⁷ Huck is serious about keeping the promise – he is not tempted to break it even when he later learns that there is a significant reward for anyone who would turn Jim in. A second incident comes a little later in the story. Huck plays a cruel prank on Jim and Jim is deeply hurt. Upon seeing Jim's disappointment, Huck finds himself apologizing to Jim: "It was fifteen minutes before I could work myself up to go and humble myself to a nigger – but I done it, and I warn't ever sorry for it afterwards, neither. I didn't do him no more mean tricks, and I wouldn't done that one if I'd a knowed it would make him feel that way."⁸⁸

The fact that this isn't just an isolated instance of an interaction with Jim in which Huck does the right thing is important for our assessment of Huck's action. It gives us good reason to believe that when Huck lies to protect Jim, it's not just a fluke that he has performed the right action. We have good reason to think that he wouldn't have given Jim away, had Huck been offered some more money in return. We also have good reason to believe that he wouldn't have given Jim away if he had been bored or slightly annoyed with Jim. This means that we have good reason to believe that Huck wouldn't have been led to do the wrong thing by the presence of certain morally irrelevant features. We also have good reason to believe that he was sensitive to some of the features that are in fact morally relevant. For example, the story gives us good reason to believe that Huck would not have lied to protect Jim from discovery if this weren't in Jim's best

⁸⁷ Twain (1884), Chapter VIII.

⁸⁸ *ibid*, Chapter XV.

interest. Huck wasn't motivated, for example, by a blind obsession with Jim's freedom – he would not have prevented Jim from returning home if Jim had decided that he really needed to go back to his family.

It's crucial for our attribution for moral worth that Huck's performance of the right action is counterfactually robust in these ways – that he wouldn't have been derailed from doing the right thing by certain morally irrelevant features and he would have adjusted his course of action in response to certain morally relevant features. Thus Huck's action is morally praiseworthy if and only if he has the ability to perform what is in fact the right action not only under the actual circumstances but also under some range of different circumstances. But this is just to say that Huck's action is only morally praiseworthy if and only if he exercised his practical moral knowledge in performing it. The counterfactual robustness that we require for Huck's action to be morally praiseworthy is exactly the same counterfactual robustness that is exhibited when an agent has practical moral knowledge. When an agent has practical knowledge how to perform some action, she has the ability to successfully perform it not only under the actual circumstances but under a range of conditions. Moral knowledge then is crucial to our evaluation of Huck's action as morally praiseworthy after all: we regard Huck's action as morally praiseworthy *because* we take him to be acting from practical moral knowledge.

4.2 How did Huck acquire Practical Moral Knowledge?

But is it plausible to think that Huck really does have practical moral knowledge? There are a few remaining worries that you might have. For one, you might think it's implausible that practical and theoretical knowledge could come apart to the degree that they would have to if Huck really does have practical moral knowledge. After all, Huck most certainly lacks theoretical moral knowledge, not just about whether it's permissible for him to lie in this particular case but also on a whole range of moral questions about slavery and the moral relevance of race. To address this worry, it will be helpful to first pin down what this practical moral knowledge that Huck has amounts to. What exactly is he able to do? There is a more general question lurking in the background here about what exactly the abilities in practical moral knowledge are. This is a somewhat tricky question: those abilities aren't as neatly identified and categorized as the ability to ride a bike or to perform CPR on a patient. In Huck's case there are a couple of obvious candidates: one is that Huck has the ability to treat *Jim* fairly. I think that this isn't giving Huck

enough credit – we’ll see why in a moment. A more plausible candidate is that Huck has some ability to treat others fairly regardless of their race. The qualification “some ability” is crucial here, since Huck obviously is not a moral saint and he isn’t entirely free from the biases that are all around him. But this shouldn’t trouble us too much: practical knowledge in other domains comes in degrees, so there’s nothing surprising about practical *moral* knowledge coming in degrees also.

If Huck’s practical moral knowledge consists in his ability to treat others fairly regardless of their race, the question is whether someone could have that ability while lacking the relevant theoretical moral knowledge – while falsely believing that race is a morally relevant feature. This is, at least in part, an empirical question. And the evidence suggests that for many nonmoral domains the answer to this question is that yes, an agent’s ability and her theoretical knowledge can diverge quite significantly. For some nonmoral cases this is not at all surprising. An expert athlete may make a terrible coach: it’s quite possible to be an excellent skier without having any insight into how one adjusts one’s movements to the slope and the snow conditions. In fact, it’s possible to be an excellent skier while having false beliefs about the correct way to ski. But, perhaps more surprisingly, this kind of opacity also holds for mental activities including problem-solving or making evaluative judgments. There is, for example, good evidence that a physicist has no more insight into how she solves a physics problem and a doctor no more access to how she arrives at a diagnosis than the skier to how she moves her muscles and adjusts her balance.⁸⁹ Doctors, for example, are experts at making clinical judgments. Yet they turn out to have false beliefs about how they make these judgments: while their judgments are reliable, they are mistaken about which method they use to arrive at these judgments. Note that in this case, it’s not just self-knowledge that the doctors lack. Since the actual method they unbeknownst to themselves employ is reliable, it responds to the medical features of the situation that are in fact medically relevant. The method they think they employ, on the other hand, is unreliable; it fails to respond to some of the medical features of the situation or it is mistaken about which weight they should be given. This means that the agents are not just mistaken about which method they

⁸⁹ Wallis [2008], p. 130 reports:

Nisbett & Wilson (1977) summarize their extensive review of the psychological literature by saying that with regard to high-level reasoning (such as that involved in problem-solving, making evaluations, making judgements, and initiating actions) subjects; (1) often had no awareness of environmental factors or information shown by analysis to have importantly influenced their responses, (2) often remained unaware of their responses, and (3) were unaware of any connection in their thought between such environmental/problem factors and their responses.

employ but also about which features of the situation are medically relevant and how much weight they should be given. Thus, the agents have the ability to make reliable medical judgments while having false medical beliefs – beliefs about which features are medically relevant.⁹⁰ Note that the divergence in this case looks analogous to the one we find in the case of Huck Finn: Huck has the ability to treat others fairly regardless of their race while having false beliefs about the moral relevance of race. If such divergence does frequently arise in nonmoral cases, it seems plausible to think that it could equally arise in the moral case.

Still, you might have a second worry about Huck's case. How could Huck have possibly acquired the practical moral knowledge in question? If growing up in a racist society means that he is not in an epistemic position to have theoretical moral knowledge, how could it possibly be compatible with his acquiring practical moral knowledge? We have noted above that something like stereotyping plays an important role in practical knowledge. It allows an agent to recognize situations that are morally equivalent *as* morally equivalent; it allows an agent to recognize situations in which someone is being water-boarded to get him to testify as morally equivalent to one in which someone is being beaten to get him to testify: both are instances of torture. This means that to have practical moral knowledge an agent must possess certain stereotypes. But something more is required. Note that to recognize morally equivalent situations as morally equivalent, the agent must not be distracted by the presence of some morally *irrelevant* factors. And this, in turn, requires the absence of some pernicious stereotypes – stereotypes that would lead the agent to see morally equivalent situations as morally different. Thus, the agent must not have stereotypes that would lead her to see a case of water-boarding, for example, a Muslim and beating a Christian as morally different. Practical moral knowledge involves both being sensitive to morally relevant features and not being sensitive to morally irrelevant features. Stereotyping plays a role in both. Thus, practical moral knowledge involves having the right kinds of stereotypes and lacking the wrong kinds of stereotypes.

This suggests that what we need to explain is why, if Huck isn't spared from the racist beliefs of his society, he nevertheless does not seem to have internalized the stereotypes that lead his contemporaries to act as if harming a white person as having very different moral status from harming a black person. We can identify two relevant factors: for one, recall Huck's upbringing.

⁹⁰ Wallis [2008], p. 131.

Huck grows up very much a social outcast: his father is a drunkard; he doesn't go to school but is free to roam around and spend his time with whoever he wants to associate with and with whoever will want to associate with him. And importantly, he spends a lot of time in the company of slaves – in the role of companions rather than slaves. In fact, slaves are often the only ones who will associate with him (and, occasionally, feed him). “Respectable” adults view Huck with suspicion; he's considered bad company for their own children. One consequence of Huck's upbringing is that he is not subject to the moral training that his peers, who grow up in more orderly households, go through. His peers don't just observe that their society is structured by race. They come to actively participate in these structures. They constantly encounter situations in which they or their families interact with slaves and they learn “the right way” to respond to these situations. They don't just observe others make moral judgments about what is a fair treatment of a slave but they also make these judgments themselves and have them corrected. So one factor is that while Huck's outcast role hasn't sheltered him from the prevalent racist beliefs, it does mean that he was less immersed in some of its social practices.

Even more important, however, is Huck's friendship with Jim. In fact, Huck and Jim's relationship is a typical example of the kinds of interracial interactions that have been shown to reduce prejudice. According to psychological research that dates back to the 1950s, an important mechanism for changing stereotypes is social contact under certain favorable conditions; this is the “social contact hypothesis”:

Since the 1950s when the social contact hypothesis was first proposed, social psychologists have distilled the conditions that contribute to a debiasing environment. People must be: (1) exposed to disconfirming data; (2) interact with others of equal status; (3) cooperate; (4) engage in non-superficial contact; and (5) receive clear norms in favor of equality.⁹¹

More recent studies lend support to the hypothesis and moreover suggest that none of these conditions is *necessary* for social interactions to successfully reduce stereotypes. We can see that the relationship between Huck and Jim clearly fits at least three of these conditions: Huck and Jim spend a lot of time together; they get to know each other very well. Their contact is not superficial. They are engaged in co-operative activity: they share the goals of finding shelter, finding ways of keeping the raft afloat, avoiding discovery. And they engage in this activity on equal footing. There is also some reason to think that in the course of his friendship with Jim,

⁹¹ Kang & Banaji [2006], p.1101.

Huck is exposed to some disconfirming evidence about members of the other race: at one point in the story, Huck muses about how Jim misses his family and how blacks care about their loved ones in the same way as white people do.⁹²

Two things are important here: One is that when social contact results in psychological change – when it reduces agents’ stereotypes – it doesn’t do so by leading the agent to revise their explicit beliefs. In fact, often such interactions will change the agent’s behavior towards members of the other group while leaving her explicit beliefs about race intact.⁹³ Thus, it’s plausible to think that his friendship with Jim led Huck to acquire practical moral knowledge without his acquiring theoretical moral knowledge. Secondly, you might be tempted to conclude that as the friendship between Huck and Jim developed, Huck did treat Jim fairly just because Jim was his friend. In other words, you might think that Huck wasn’t really responding to morally relevant features, but merely to the friendship. Thus, you might think Huck’s action was only accidentally right: he happened to do the right thing just because Jim happened to be his friend. But this would be a mistake. What is true is that the friendship with Jim was a necessary condition for Huck to acquire a certain moral ability: the ability to treat others fairly regardless of their race. And it may be true that Huck was lucky to have had the chance to befriend Jim. But once Huck did acquire the practical moral knowledge in question, it wasn’t a matter of luck that he did the right thing in the situation he found himself in.

Huck’s action seems morally praiseworthy. But Huck does not have theoretical moral knowledge. So, theoretical moral knowledge cannot be necessary for moral worth. However, I have argued, Huck does have practical moral knowledge – his ability to do the right thing gives us a plausible explanation for why his action is morally praiseworthy. Huck’s case then suggests that moral knowledge is necessary for moral worth after all. And it suggests that practical moral knowledge can be sufficient for moral worth, even in absence of theoretical moral knowledge.

6. An Alternative Account

⁹² For more discussion of what kind of interactions with a minority lead to a reduction of implicit biases, see Anderson [2010].

⁹³ Kang and Banaji [2006] report that while the number of friends of the other race affects an agent’s implicit attitudes, “the friendship measure had no correlation to measures of explicit bias.” See p. 1103.

On the view I am defending, moral knowledge is a necessary condition for moral worth. Acting from one's moral knowledge is also sufficient for moral worth. The moral knowledge in question can be theoretical moral knowledge. To act from one's theoretical moral knowledge an agent needs to also be motivated by her (noninstrumental) desire to do the right thing. My account therefore accommodates intuitively praiseworthy actions like John's in Whistle-Blowing. But moral knowledge can also be practical and in this case it can guide an agent's actions directly, without the agent having to make an explicit moral judgment first. I have argued that this accounts for the moral worth of Huck Finn's action. What I want to do in this section is to compare my account to a rival one by Arpaly. Arpaly's proposal, I will argue, does not give us a satisfactory account of the counterfactual robustness of morally praiseworthy actions.

6.1 Acting on the right kind of desires

On this rival account moral worth depends solely on the *desires* motivating the agent. Arpaly's account is particularly motivated by Huck Finn's case. Arpaly argues that since Huck Finn does not believe that protecting Jim is the right thing to do, his moral beliefs must be irrelevant to whether an action is morally praiseworthy:

Usually, "acting from duty" and "acting for the sake of the fine" are taken to indicate not simply acting for moral reasons, but rather acting for reasons believed or known to be moral reasons—acting out of a desire to do that which is right. But moral worth is fundamentally about acting for moral reasons, not about acting for reasons believed or known to be such, and distinguishing the two is important in evaluating moral agents.⁹⁴

Arpaly attributes the moral worth of Huck's action not to his moral knowledge but rather to his having the right kinds of desires. She argues that whether an agent acts from moral reasons depends on her motivating desires.

An action is morally praiseworthy if and only if it is motivated by moral concern:

For an agent to be morally praiseworthy for doing the right thing is for her to have done the right thing for the relevant moral reasons—that is, for the reasons for which the action is right (the right reasons clause); and an agent is more praiseworthy, other things being equal, the deeper the moral concern that has led to her action (the concern clause). Moral concern is to be understood as concern for what is in fact morally relevant and not as concern for what the agent takes to be morally relevant.⁹⁵

⁹⁴ Arpaly [2003], p. 73.

⁹⁵ *ibid.*, p. 84.

And to have moral concern is to have desires for the features of the situation that are *in fact* morally relevant – whether or not the agent believes or knows them to be relevant:

To say that a person acts out of moral concern is to say that a person acts out of an intrinsic (noninstrumental) desire to follow (that which in fact is) morality, or a noninstrumental desire to take the course of action that has those features that make actions morally right...⁹⁶

I have argued that Huck is praiseworthy for his action because he acts from practical moral knowledge. Arpaly proposes an alternative explanation: Huck is morally praiseworthy because, despite his false beliefs, Huck's desires track the features of the situation that are morally relevant. Huck is moved to act by his desire for Jim's freedom or perhaps Jim's well-being. Jim's freedom and well-being are morally relevant considerations even though Huck fails to believe that they are. Thus Huck's action is motivated by his moral concern – a desire – for a right-making reason – Jim's freedom.

6.2 Desires and Counterfactual Robustness

I do think that there may well be something morally attractive about an agent who has certain kinds of desires (for example, the desire for others to be happy) and, perhaps more importantly, lacks other kinds of desires (the desire to see others suffer). But there are reasons to be skeptical that moral worth could be a matter of just acting on the right kind of desires. We have seen that it's crucial for our taking an action to have moral worth that the action be counterfactually robust in the right kind of way: that it isn't just an accident that the agent performed the right action but that she would have done the right thing even if conditions had been different. The aim of an account of moral worth is to identify the feature of the agent in virtue of which this counterfactual robustness obtains. Desire is not a plausible candidate for this feature. When an agent is motivated by the right kind of desire but lacks moral knowledge, her action isn't counterfactually robust in the required way.

To see that something is amiss in Arpaly's account, consider her "concern clause". Arpaly argues that an action has greater moral worth when it's the result of deeper moral concern, i.e. a stronger desire for the morally relevant feature. The motivation for this clause is the intuition that an agent seems to deserve more moral credit when her doing the right thing is particularly

⁹⁶ *ibid.*

counterfactually robust – when she would have performed the action even under particularly difficult conditions. Take, for example, two agents who are both doing well and who both make a donation to charity. One of them, the “die-hard philanthropist”, gives money to charity even when things aren’t going well for him, while the “fair-weather philanthropist” only gives money to charity occasionally, when she happens to be in a particularly good mood. There’s an intuition that the “die-hard philanthropist” deserves more moral credit for his action, even when as a matter of fact both philanthropists are motivated by the fact that their donation will help the needy. After all, it seems that the “die-hard philanthropist” would have done the right thing even if he weren’t in a good mood and even if things were going poorly for him. His right action has greater counterfactual robustness than that of the “fair-weather” philanthropist.

Arpaly argues that this intuition is right: the reason the “die-hard philanthropist” deserves more moral credit for giving to charity even when things are going well for him is that his desire to help the needy is *stronger* than that of the “fair-weather philanthropist”. The strength of the desire explains the greater counterfactual robustness. But this seems odd: there isn’t a close connection between strength of desire and counterfactual stability. There’s nothing paradoxical about saying that someone has an extremely strong desire that’s very fickle. I may, for example, experience an overwhelming desire to drink some coffee when and only when I can smell it being freshly ground. Once the smell is gone, so is the desire. Similarly, we can imagine that the “fair-weather philanthropist”, when she’s feeling good, is consumed by an extremely strong desire to help others: she can think of and do nothing else. Her actions can be motivated by a desire that is both: extremely strong and very fickle. Strength of desire simply doesn’t track which actions are more or less morally praiseworthy.

But it’s not just that desires cannot account for why some actions may seem more praiseworthy than others. Desires cannot give us the right counterfactuals at all. Consider again the case of Huck Finn: we think that Huck is praiseworthy because it’s not just an accident that he did what was right. Is there a desire that can explain why Huck’s action is non-accidentally right in the relevant way? One obvious candidate is a desire for Jim’s freedom. But this cannot do: Huck wouldn’t be morally praiseworthy if he had insisted on freeing Jim, even if Jim absolutely didn’t want to be freed – if he wanted to stay with Miss Watson to be near his wife and children. Could it be a desire for Jim’s well-being? Again, it’s not clear that a mere desire for Jim’s well-being will

guarantee that Huck will do the right thing. After all, we wouldn't think that Huck was morally praiseworthy if he acted in Jim's interest *no matter what* – for example, if Jim were a known serial killer hiding from the authorities and gearing up for his next killing spree.

My criticism here echoes the Kantian arguments for why mere inclination is not enough for morally worthy actions: we need more than just the right kinds of desires. Even if an agent has desires for all kinds of morally relevant features and these various desires lead her to perform the right action in a given situation, the balance of the desires seems rather precarious. A morally relevant feature that weighs little in one situation may weigh a lot in a different one. When it does, the agent's set of desires will lead her morally astray. As Herman argues in criticizing what she calls the Humean picture (which turns out to be very close to the one advocated by Arpaly):

If we begin with the distinct motives that mark the different regions of moral attention and concern..., it does not follow that the cluster of motives judged to be good can co-exist in one person. We want persons to be moved by considerations of gratitude, humanity, natural affection, generosity, and industry; but there is no natural, inevitable fit, no general template, that directs their joint instantiation.⁹⁷

For an action to have moral worth, the agent needs to be sensitive to the moral status of the action – to whether it's right or wrong. A desire for a morally relevant feature is not enough to lead the agent to discriminate between right and wrong actions. After all some morally relevant feature can be outweighed by others that are present – an agent can hence do the wrong thing while still being motivated by a feature that is morally relevant. Contrary to what Arpaly claims, desires for morally relevant features simply are not the same as *moral concern*. This problem does not arise for my proposal. This is because for an agent to have moral knowledge, it's not enough that the agent be sensitive to some morally relevant *feature*. Moral knowledge requires the agent to be sensitive to the *moral status* of the action: the agent must either have theoretical knowledge that the action is right or she must have the ability to perform the right action under a range of conditions.

7. Conclusion

I have argued that moral knowledge is central to moral worth: whether an agent's action is morally praiseworthy depends on whether she is acting from moral knowledge. One advantage of this account is that it allows us to accommodate a wide range of intuitively praiseworthy actions.

⁹⁷ Herman [2007], p. 86.

This is because moral knowledge includes both practical and theoretical moral knowledge. Thus, agents can be morally praiseworthy when they perform the right action upon having reflected what they ought to do. But they can also be morally praiseworthy when, like Huck, they have the ability to do the right thing despite having false beliefs about what morality requires. Another advantage is that it accounts for the counterfactual robustness of morally praiseworthy actions – the fact that it wasn't just an accident that the agent did what morality required.

Both practical and theoretical moral knowledge are attractive features of an agent's character. Ideally, we want agents to have both kinds of moral knowledge. Much like being an expert surgeon involves knowledge of anatomy and surgical procedures as well as the ability to perform surgery, so being an ideal moral agent involves knowing what morality requires and having the ability to do the right thing. The account of moral worth that I have defended hence naturally fits with an account of virtue as moral expertise.⁹⁸ The feature of the agent in virtue of which her actions have moral worth – moral knowledge – is precisely the one which she shares with the moral ideal. At the same time, it is compatible with the fact that morally imperfect agents can perform morally praiseworthy actions: Huck's practical moral knowledge allows him to compensate for his false moral beliefs. And in a novel situation, an agent can rely on moral advice to compensate for her lack of ability to do the right thing.

Finally, the account of moral worth I have sketched provides us with a definite answer to the question why we should care about moral worth at all – when an action is right, why should it matter whether it's also morally worthy or not? The answer is that moral worth matters because rightness matters and we should care about moral worth because we care about agents doing the right thing: You can't do the right thing if you don't know what it is.

⁹⁸ For an example of such an account, see Hagen [ms].

Bibliography

- Adams, Marcus P. (2009). Empirical evidence and the knowledge-that/knowledge-how distinction. *Synthese* 170 (1):97 - 114.
- Anderson, Elizabeth (2010). *The Imperative of Integration*. Princeton University Press.
- Anscombe, Elizabeth (1981), Authority in Morals. In: *The Collected Philosophical Papers of G.E.M. Anscombe, iii: Ethics, Religion, and Politics*, Oxford University Press.
- Annas, Julia (2003), The Structure of Virtue. In: DePaul, Michael and Zagzebski, Linda (eds.) *Intellectual Virtue: Perspectives from Ethics and Epistemology*. Oxford University Press.
- Arpaly, Nomy (2003). *Unprincipled Virtue: An Inquiry Into Moral Agency*. Oxford University Press.
- Bargh, John and Chartrand, Tanya (1999). The Unbearable Automaticity of Being. *American Psychologist* 54 (7):462-479.
- Barron, Sean and Grandin, Temple (2005). *The Unwritten Rules of Social Relationships*. Future Horizons.
- Baron, Marcia (1984). The alleged moral repugnance of acting from duty. *Journal of Philosophy* 81 (4):197-220.
- Bennett, Jonathan (1974). The Conscience of Huckleberry Finn. *Philosophy* 49 (188):123-134.
- Christensen, David (2007). Epistemology of disagreement: The good news. *Philosophical Review* 116 (2): 187-217.
- Christensen, David (2010). Higher order evidence. *Philosophy and Phenomenological Research* (81) 1: 185-215.
- Coady, C. A. J. (1992). *Testimony: A Philosophical Study*. Oxford University Press.
- Driver, Julia (2001). *Uneasy Virtue*. Cambridge University Press.
- Driver, Julia (2006). Autonomy and the asymmetry problem for moral expertise. *Philosophical Studies* 128 (3): 619 - 644.
- Fricker, Miranda (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford University Press.
- Gibbard, Alan, (1990), *Apt Choices, Wise Feelings*. Harvard University Press.
- Glick, Ephraim (2011). Two Methodologies for Evaluating Intellectualism. *Philosophy and Phenomenological Research* 83 (2):398-434.

- Glick, Ephraim (2009). Practical Knowledge and Abilities. Ph.D. dissertation, MIT.
- Hagen, Daniel (ms.). Virtue as Expert Moral Knowledge: A Socratic Account. Unpublished manuscript.
- Harman, Elizabeth (2011). Does moral ignorance exculpate? *Ratio* 24 (4):443-468.
- Haslanger, Sally (2007). "But mom, crop-tops are cute!" Social knowledge, social structure and ideology critique. *Philosophical Issues* 17 (1):70-91.
- Haslanger, Sally (2010). Ideology, Generics, and Common Ground. In Charlotte Witt, ed., *Feminism and Metaphysics*. Springer.
- Hawley, Katherine (2003). Success and Knowledge-How. *American Philosophical Quarterly* 40 (1):19 - 31.
- Hawley, Katherine (2010). Testimony and knowing how. *Studies in History and Philosophy of Science Part A* 41 (4):397-404.
- Henson, Richard G. (1979). What Kant might have said: Moral worth and the overdetermination of dutiful action. *Philosophical Review* 88 (1):39-54.
- Herman, Barbara (1981). On the value of acting from the motive of duty. *Philosophical Review* 90 (3):359-382.
- Herman, Barbara (2007). *Moral Literacy*. Harvard University Press.
- Hills, Allison (2009), Moral Testimony and Moral Epistemology. *Ethics* 120 (1): 94-127.
- Holton, Richard (2009). *Willing, Wanting, Waiting*. Oxford University Press.
- Holton, Richard (ms). Inverse akrasia and weakness of will.
- Hopkins, John (2007), What is Wrong with Moral Testimony? *Philosophy and Phenomenological Research* 74 (3): 611-634.
- Jones, Karen (1999), Second-Hand Moral Knowledge. *The Journal of Philosophy* 96 (2): 55-78.
- Jones, Karen (2003). Emotion, Weakness of Will, and the Normative Conception of Agency. In A. Hatzimoysis (ed.), *Philosophy and the Emotions*. Cambridge University Press.
- Lackey, Jennifer (2008), *Learning From Words: Testimony as a Source of Knowledge*. Oxford University Press.
- Lillehammer, Hallvard (1997). Smith on moral fetishism. *Analysis* 57 (3):187-195.

- Kang, J. & Banaji, M. R. (2006). Fair Measures: A Behavioral Realist Revision of 'Affirmative Action'. *California Law Review*, Vol. 94, No. 4: 1063-1118.
- Kant, Immanuel (1785/1996), *Groundwork of the Metaphysics of Morals*, in: Gregor, Mary (trans. and ed.) *Practical Philosophy*. Cambridge University Press.
- Kearns, Stephen & Star, Daniel (2009). Reasons as Evidence. *Oxford Studies in Metaethics* 4:215-42.
- Kennett, Jeanette (2002). Autism, empathy and moral agency. *Philosophical Quarterly* 52 (208):340-357.
- Korsgaard, Christine M. (2008). *The Constitution of Agency: Essays on Practical Reason and Moral Psychology*. Oxford University Press.
- Markovits, Julia (forthcoming), Saints, Heroes, Sages, and Villains.
- Markovits, Julia (2010), Acting for the Right Reasons. *Philosophical Review* 119 (2): 201-242.
- McGrath, Sarah (2011), Skepticism about Moral Expertise as a Puzzle for Moral Realism. *Journal of Philosophy* 108 (3).
- McGrath, Sarah (2009). The puzzle of pure moral deference. *Philosophical Perspectives* 23 (1): 321-344.
- Nickel, Philip (2001), "Moral Testimony and Its Authority", *Ethical Theory and Moral Practice* 4 (3): 253-266.
- Smith, Michael (1994). *The Moral Problem*. Blackwell.
- Schwitzgebel, Eric (2010). Acting Contrary to Our (Professed) Beliefs. *Pacific Philosophical Quarterly* 91:531---553.
- Stanley, Jason (2011a). Knowing (How). *Noûs* 45 (2):207-238.
- Stanley, Jason (2011b). *Know How*. Oxford University Press.
- Stanley, Jason & Williamson, Timothy (2001). Knowing how. *Journal of Philosophy* 98 (8): 411-444.
- Stocker, Michael (1976). The schizophrenia of modern ethical theories. *Journal of Philosophy* 73 (14):453-466.

- Twain, Mark (1884). *Huckleberry Finn*. Edited by Walter Blair and Victor Fisher (1988), University of California Press.
- Wallis, Charles (2008). Consciousness, context, and know-how. *Synthese* 160 (1):123 - 153.
- Williams, Bernard (1981). Persons, Character, and Morality. In *Moral Luck*. Cambridge University Press.
- Wolff, Robert Paul (1990), The Conflict Between Authority and Autonomy. In Raz, J. (ed.) *Authority*, Blackwell.