

MORAL EXPERTISE, MARGINALISATION, & DEFERENCE

Grace Boey

A DISSERTATION

in

Philosophy

Presented to the Faculties of the University of Pennsylvania

in

Partial Fulfilment of the Requirements for the

Degree of Doctor of Philosophy

2020

Supervisor of Dissertation

Graduate Group Chairperson

Errol Lord

Errol Lord

Errol Lord

Errol Lord

Associate Professor of Philosophy

Associate Professor of Philosophy

Dissertation Committee

Errol Lord, Lisa Miracchi, Daniel Singer

Dedication

For Ada and the ducks.

Acknowledgements

Thank you to Errol Lord, Lisa Miracchi, and Daniel Singer, for guiding me intellectually, personally, and professionally.

For suggesting the topic for a paper that would eventually grow into this dissertation, I thank Alex Guerrero.

Thank you to the members of the Penn Normative Philosophy Group, and the Mind-Like, Intelligence, and Research Group for helping me develop these ideas.

For being the best cohort mates one could ask for, I thank Eilidh Beaton, Max Lewis, Michael Vazquez, and Zhang Mingjun.

Thank you to Jeff Botula for being there throughout.

For supporting me through this doctorate, thank you to Kay Adjoa-Oko, Ben Baker, Conrad Bakka, Jenna DiLossi, Dylan Manson, Wanda Lee, Reggie Middlesdorf, Lisa Seufert, Chirsell Tah, Sean Thomas, and Shaun Wellesley.

This doctorate was made possible by Nanyang Technological University's HASS International PhD Scholarship. I thank past and present NTU colleagues and friends, including Alan Chan, Li Chenyang, Andrés Luco, Christina Chuang, Preston Greene, Lina Jansson, Teru Miyake, and Winnie Sung.

Abstract

MORAL EXPERTISE, MARGINALISATION, & DEFERENCE

Grace Boey

Supervisors: Errol Lord, Lisa Miracchi, Daniel Singer

In this dissertation, I make the broad case against overestimating the moral epistemic capacities of adult moral agents, and the demands placed upon them to figure things out for themselves, morally speaking.

This work is split into three chapters. In Chapter 1, 'Moral Expertise & Experience', I argue that moral expertise in some moral sub-domain, or one's competence at forming moral knowledge in response to morally relevant features within that moral sub-domain, is typically generated through experience with the concrete world. I reject the claim that moral philosophers are the best candidates for being moral experts, and that imagination provides us an equally good path towards moral expertise as experiences does.

In Chapter 2, 'Moral Expertise on Oppression', I argue that oppressed group members are often in a better position to become a moral expert with respect to the type of oppression experienced by that group. For instance, women are in a better position to become moral experts with respect to the moral sub-domain of sexism, as opposed to men.

In Chapter 3, 'Virtuous & Worthy Moral Deference', I vindicate the practice of deference to second-hand moral testimony, and agents who defer, in the face of what I call the 'reasons unresponsiveness observation'. This is because in certain contexts, it's important for them to be motivated by a concern for doing the right thing in itself. Pure moral deference, practiced in the right way, can allow us to exercise virtues, and act with moral worth.

Table of Contents

	Introduction	1
1	Moral Expertise & Experience	7
1.1	Introduction	7
1.2	Moral expertise & moral experts	11
1.3	The non-intellectualist picture of moral expertise	19
1.4	The role of experience in producing moral experts	28
1.5	The challenges of imagination	48
1.6	Conclusion	70
2	Moral Expertise on Oppression	72
2.1	Introduction	72
2.2	Moral expertise	75
2.3	Moral expertise on oppression	86
2.4	Conclusion	101
3	Virtuous & Worthy Moral Deference	105
3.1	Introduction	105
3.2	The threat of reasons unresponsiveness	110
3.3	The contextual claim about moral worth	118
3.4	Virtues of acting & knowing for rightness' sake	129
3.5	Conclusion	138
	Bibliography	144

Introduction

Many of us strive to better ourselves as moral agents. This process begins at childhood, where we learn by example and instruction from authority figures. At this stage, we are often told what's right and wrong, and learn by accepting the moral testimony of our parents and teachers. As we progress through childhood and adolescence, we typically become better at figuring out what's right and wrong on our own, and learn how to successfully navigate increasingly complex situations. And once we reach adulthood, it's commonly thought that we become fully fledged moral agents.

A common intuition is that once we become adults, the onus is then on us to figure things out by ourselves, morally speaking. It becomes less palatable, or even taboo, to simply take someone else's word about what's right or wrong, as we did when we were children. For example, consider the following case:

Vegetarian. Eleanor has always enjoyed eating meat but has recently realised that it raises some moral issues. She knows lots of empirical facts about animals and farming. But rather than thinking further about these, 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. (Adapted from Hills 2009, 94)

Assuming that eating meat is indeed wrong, it looks like Eleanor's friend is more morally competent than her at forming moral knowledge - at least with respect to the matter of eating meat, and perhaps animal ethics in general. Still, assuming that Eleanor is an adult, many will find something suspicious, or undesirable, about what she's done. As an adult moral agent, the onus is on Eleanor to think through these issues for herself.

Eleanor's engaged in the epistemic practice of what's been called *pure moral deference*. This is the practice of forming and sustaining some moral belief solely based on someone else's purely moral testimony. What makes this 'pure', rather than 'impure', moral deference is that the testimony and deference is of an entirely moral, non-empirical nature. Assume that Eleanor knows all the relevant empirical facts that make eating meat wrong, including facts about animal suffering, the factory farming system, and so on. What she's lacking is competence at putting these pieces of the moral puzzle together. So, the new belief she accepts is purely moral, and isn't formed or sustained on the basis of any empirical information she lacks.

The *Vegetarian* case might make one think that pure moral deference is impermissible or undesirable. But the intuition becomes less compelling once we think about cases like these two:

Halloween. Joel is attending an Intro to Ethics recitation. The class is discussing the ethics of Halloween costumes, which he's never thought about before.

Because Joel has a history midterm right after class, he isn't paying attention. Rather, he's reciting historical facts about minstrel shows in his head. At the very end of class, Joel hears his TA state that her own view is that donning blackface for Halloween is wrong. Joel leaves class having no clue why his anyone would hold this view. He also has no feelings about either blackface or minstrel. He knows they upset the African American community, but thinks everything seems to upset someone anyway. Still, he believes his TA, and trusts her too. Whenever she's explained her moral views in the past, he's always been able to grasp her explanations, and has always agreed with her conclusions. Joel, a white man, decides not to paint his face black and dress up as Black Panther for Halloween. When asked why, Joel sighs and says, 'It would be wrong. I'm not sure why, but my TA said so.'

Ableism. Kim is an able-bodied person. Alex, her good friend, uses a wheelchair as both his legs are permanently injured. Kim knows that ableism is wrong, is committed to eliminating it, and is able to identify some obvious instances of it on her own. But on numerous occasions over their twenty year friendship, she's had to rely on Alex to explain other instances of ableism to her.

Fortunately, Kim has always grasped Alex's patient explanations, and has always come to agree with his conclusions. This process has sometimes taken her many days, or even weeks. But she's always eventually seen the light,

even if she's initially disagreed. Kim has also come to know lots of descriptive facts about how Alex has to navigate the world. Less fortunately, Kim still hasn't managed to become much better at identifying ableism on her own.

One day, Alex complains to Kim that his new boss has started pushing him around, without asking him if he needs help. Kim doesn't see how this is ableist. Alex tries to explain how this undermines his bodily autonomy, but after five minutes, Kim still doesn't get it. She refuses to believe Alex.

Although something seems off about Joel in *Halloween*, there still seems to be something good about what he's done. As for Kim in *Ableism*, she seems to have strong reason to take Alex's moral testimony that pushing him around without his permission is ableist. In fact, we might think that Kim is required to do so, and that it's morally and epistemically wrong for her to refuse.

The broad aim of this dissertation is to show how cases like *Halloween* are good, and how there is something desirable about what Joel has done. It also aims to motivate the thought that agents similar to Alex and Kim are quite common, where one's moral sensitivities are heightened or blunted depending on one's social status and experiences with being marginalised. Ultimately, the dissertation makes the case against overestimating the moral epistemic capacities of adult moral agents, and the demands placed upon them to figure things out for themselves.

The dissertation is split into three chapters. Chapter 1, 'Moral Expertise & Experience', argues that moral expertise is typically generated through experience with the concrete world. I argue that moral expertise in some moral sub-domain, or one's competence at forming moral knowledge in response to morally relevant features within that moral sub-domain, is typically generated through experience with the concrete world. I reject the claim that moral philosophers are the best candidates for being moral experts, and that imagination provides us an equally good path towards moral expertise as experiences does.

In Chapter 2, 'Moral Expertise on Oppression', I argue that oppressed group members are often in a better position to become a moral expert with respect to the type of oppression experienced by that group. For instance, women are in a better position to become moral experts with respect to the moral sub-domain of sexism, as opposed to men. And in *Ableism*, Alex is in a better position to become a moral expert with respect to the moral sub-domain of ableism, as opposed to Kim.

In Chapter 3, 'Virtuous & Worthy Moral Deference', I vindicate the practice of pure moral deference, and agents who defer, in the face of what I call the 'reasons unresponsiveness observation'. This is because in certain contexts, it's important for them to be motivated by a concern for doing the right thing in itself. In particular, I'll argue for the following claim:

The contextual claim about moral worth: In contexts *C* where an agent's unresponsiveness to the moral reasons why φ -ing is right makes her

incompetent at φ -ing on the basis of those moral reasons, her act of φ -ing has moral worth if and only if (i) she's motivated to φ by a concern for doing the right thing, and (ii) she's competent at φ -ing based on this concern.

I'll also argue for the existence of the following two virtues:

The virtue of acting for rightness' sake: An agent's competence at using her concern for doing the right thing to perform right actions, in contexts where because of her unresponsiveness to moral reasons, she'd otherwise be incompetent at doing so.

The virtue of knowing for rightness' sake: An agent's competence at using her concern for doing the right thing to form moral knowledge, in contexts where because of her unresponsiveness to moral reasons, she'd otherwise be incompetent at doing so.

Pure moral deference, practiced in the right way, is a way for agents to meet the standards of the contextual claim. This allows them to exercise virtues, and act with moral worth.

Chapter 1

Moral Expertise & Experience

1 Introduction

In navigating our everyday lives, we rely on others all the time. One way we do this is by relying on the *expertise* of others. We rely on experts in fields like science, math, medicine, engineering, and other fields. We do so in seeing the doctor, trusting that the bridge beneath us won't fall apart, trusting the translations in travel guidebooks, and calling the plumber when our sink malfunctions. In relying on the expertise of others, we recognise that they have skills and knowledge that we lack. We also recognise that experts can be trusted to deploy these skills and knowledge to achieve outcomes we'd unable to accomplish on our own. Reliance on the expertise of others is seen as a legitimate means of navigating the world.

Are there such things as *moral experts* and *moral expertise*? To some, the existence of experts and expertise in the moral domain might seem intuitively much less plausible than the existence of experts in non-evaluative domains like medicine.

There are several reasons why one might be skeptical of moral expertise. One is that experts are those to whom we ought to defer, and it seems strange to defer to anyone about morality. Indeed, there's a lively philosophical debate about the permissibility of practicing moral deference. However, most participants in the debate don't deny that moral expertise exists, and locate the shortcomings of moral deference elsewhere.¹ Another reason one might be skeptical is that moral truths are *a priori*. Since moral truths are in principle equally available to all of us, no one has any special claim on moral knowledge. Yet, there is no reason to think that everyone is equally well-positioned to grasp truths in *a priori* domains. We recognise, for instance, the existence of experts in mathematics (McGrath 2009; 2011a). One might deny the existence of moral experts by being an anti-realist about moral facts. If there are no objective mind-independent moral facts, then the concept of moral expertise makes no sense. I take it, however, that moral realism isn't something many of us want to give up so easily.²

Of course, these don't exhaust the reasons why one might resist the idea of moral expertise. Still, many agree that moral expertise can and does exist. I shall follow suit, and assume that it does. The aim of this chapter, then, is to answer new questions that arise once we acknowledge the existence of moral expertise. *What do*

¹ Philosophers have noted, for instance, that moral deference bypasses moral understanding (Hills 2009), virtue (Hills 2009, Howell 2014, Lord 2017), and appreciative knowledge (Lord 2017). McGrath (2008; 2009; 2011a; 2011b) argues that the problem is epistemic, because moral experts are difficult for non-experts to identify.

² It has also been argued that moral expertise is still compatible with non-realist positions like relativism (Jones & Schroeter 2018, 463-464).

moral experts look like? And how does moral expertise come about?

Relatively recent work on moral expertise by Driver, Jones, Schroeter, and others paints a combined and compelling picture with at least features.³ (1) First, moral expertise is typically generated through experience. This stands in contrast with one common image of the moral expert who deliberates and delivers moral advice from the armchair—in other words, the image of moral experts as moral philosophers. (2) Second, moral expertise typically exists locally to moral sub-domains such as friendship, justice, the ethical status of killings, and so on. Moral experts aren't typically experts about 'morality in general'. (3) Third, moral expertise isn't always action guiding. We shouldn't expect moral experts to only deliver all-things-considered judgments about what is *right* and *wrong*. Rather, we should allow for the deliverances of expert judgments about what's, for instance, *cruel* or *oppressive*, or what counts as *murder*.

I find these claims plausible and compelling. I argue, however, that what's missing from the current picture is a sufficiently detailed account of claim (1)—the claim that moral expertise is typically generated through experience. This claim is crucial to the picture, because claims (2) and (3) follow directly from it. Without a detailed story of precisely how and why particular experiences put one in an especially good position to develop moral expertise, the entire picture of moral expertise is vulnerable.

³ See Jones & Schroeter (2012; 2018), Driver (2006; 2013), and less recently, Jones (1999).

Following this, the main aim of this paper will be to justify the claim that moral expertise is typically generated through experience. But in the process of doing so, this paper will also cut across other important issues and debates. For instance, I will reject the claim that moral philosophers are the best candidates for being experts in moral judgment. I will also de-emphasise the role of necessary moral truths for moral expertise, and argue that we should require moral experts to be good at gaining contingent moral knowledge about the actual world. And although I invoke sentimentalist perceptualism in my argument for how experience is important, I reject what some sentimentalist perceptualists (and many other philosophers) have recently claimed, or at least implied. This is the claim that engaging in imagination provides an equally good path towards moral expertise as experience provides.

Here's the plan. In Section 2, I'll introduce the question of whether moral philosophers count as moral experts, disambiguate between several forms of moral expertise, and define moral expertise and moral experts in *judgment*, which are the targets of my discussion. In Section 3, I'll outline what I call the 'non-intellectualist' picture of moral expertise painted by Driver, Jones, & Schroeter. First, I'll give a brief overview of the current literature. Second, I'll argue that non-intellectualists about moral expertise should also accept that moral experts must be good at making contingent moral judgments, and that their skill at making necessary moral judgments is much less important. Third, I'll argue that non-intellectualists need to

do more to justify the role of experience.

In Section 4, I'll show how experience puts agents in good positions to acquire moral expertise, and eventually become a moral expert. In doing so, I invoke the moral epistemic view of sentimentalist perceptualism. In short, experience is good at causally prompting affective moral perceptions, causing contingent moral judgments, and causing moral adjustment over time. I propose that this plausibly gives rise to a form of moral perceptual learning. In Section 5, I consider the claim that imagination is as good a means towards become a moral expert as experience. I argue that imagination in both the form of thought experiments as well as narrative fiction are inferior as compared to concrete experience with the real world.

2 Moral expertise & moral experts

2.1 Are moral philosophers moral experts?

One appealing line of thought is that moral philosophers are good candidates for being moral experts, and that moral expertise arises from philosophical training in ethics. Several philosophers have described themselves as moral experts (e.g. see Singer 1972a; Singer & Wells 1984; Crosthwaite 1995; Føllesdal 2004). Portraying what I'll call the 'intellectualist' picture of moral expertise, Singer lists the advantages that moral philosophers have over 'the ordinary man':

'Is the ordinary man just as likely to be expert in moral matters as the moral philosopher? ... The moral philosopher does have some important advantages over the ordinary man. First, his general training as a philosopher should make him more than ordinarily competent in argument and in the detection of invalid inferences. Next, his specific experience in moral philosophy gives him an understanding of moral concepts and the logic of moral argument. ... Finally, there is the simple fact that the moral philosopher can, if he wants, think full-time about moral issues, while most other people have some occupation to pursue which interferes with such reflection. ... If we are to make moral judgments on some basis other than our unreflective intuitions, we need time, both for collecting facts and for thinking about it. Moral philosophers have, then, certain advantages which could make them, relative to those who lack these advantages, experts in matters of morals.' (Singer 1972a, 117)

There is certainly something to what Singer says. Schwitzgebel also notes that a number of philosophers have also asserted that professional training helps protect them from unconscious and unwanted biases in their domain of expertise (see e.g. Grundmann 2010; Hoffman 2010; Williamson 2011; Wright 2010; from Schwitzgebel 2012, 3).

It would be nice if all this was true. Moral philosophers had better count as

experts in *something* to do with morality, otherwise it is unclear what their enterprise counts for. Unfortunately, things are less rosy and more complicated than they seem. First, empirical research suggests—embarrassingly—that professional ethicists appear to behave no differently from non-ethicists of similar social background (Schwitzgebel & Rust 2016). Second, professional philosophers, including ethicists, appear to be no less subject to framing and order effects in moral problem-solving than academics in non-philosophy fields (Schwitzgebel & Cushman 2015; 2012). Third, reading between Singer’s lines makes it clear that he thinks being *intelligent* and having superior cognitive abilities are necessary for being a moral expert. Yet research suggests that individuals with higher cognitive abilities are prone to larger bias blind spots than average (West, Meserve & Stanovich, 2012). These last two points are especially threatening to the idea that philosophical training protects moral philosophers from unconscious biases in morality. And these findings in general suggest there must be something more to being a moral expert than having the characteristics of a moral philosopher.

2.2 Defining moral expertise in judgment

Much of this confusion disappears if we distinguish between different types of moral expertise. Driver helpfully identifies three distinct types. First is expertise in moral judgment, which involves the ability to arrive at true moral judgments better than others. Second is expertise in moral practice, which involves the ability to act

morally well more than others. Third is expertise in moral analysis, which involves superior insight into the nature of morality in some respect (Driver 2013, 283). Jones & Schroeter identify a fourth type, which is expertise in giving and evaluating moral reasons (or, 'expertise in moral reasoning'). This involves the ability to articulate morally relevant features of situations (Jones & Schroeter 2018, 460).

Moral philosophers are good candidates for having expertise in moral analysis and reasoning. For the purposes of my discussion, though, I'm not interested in these types of moral expertise. Neither am I interested in expertise in moral practice, which moral philosophers (unfortunately) do not seem to excel in. What I'm exclusively interested in is a version of Driver's expertise in moral *judgment*.

Expertise in moral judgment is distinct from the other types of moral expertise in the following ways. First, it is clearly distinct from expertise in moral action; one may be good at making correct moral judgments while bad at acting on them. The distinction between expertise in moral judgment and expertise in moral analysis and reasoning is trickier. On one hand, expertise in moral judgment is conceptually distinct from the latter two types of expertise. On the other hand, one might think that having expertise in moral analysis and reasoning puts one in a good position to also excel at making correct moral judgments. Perhaps it is the case that experts in moral analysis and reasoning are highly skilled at '[arranging] moral principles in a hierarchy and so be able to deduce what ought to be done in any

given situation simply by applying the rules' (Jones & Schroeter 2012; 2018). But if Schwitzgebel & Cushman are correct that ethicists are equally subject to framing and order effects in moral problem-solving, the link between having expertise in moral analysis and reasoning and having expertise in moral judgments is not so tight after all.

Schwitzgebel & Cushman's findings are intriguing, but it would be unwise to rely too much on the scope of their preliminary studies to conclude that moral philosophers aren't especially well-placed to have expertise in moral judgment. But I will argue that there's another reason to think this is true. In short, an important component of moral expertise in judgment is the ability to make judgments about contingent moral truths in complicated situations. The methodology of moral philosophers is not a very good means towards cultivating this ability to its fullest extent. I will bracket this issue for now, and fully address it in Section 4.

So, as I've noted and as I'll show, expertise in moral analysis and reasoning doesn't put one in an especially good position to acquire expertise in moral judgment. On the flip side, one may also have expertise in moral judgment without having expertise in moral analysis or reasoning. Familiarity with moral theory doesn't seem necessary for moral knowledge. Someone could know, or justifiably believe, that lying is morally wrong without inferring this from some higher-order principle (Tropman 2018, 474). And returning to the matter of contingency, philosophers have also noted that a significant amount of our moral knowledge is of

contingent moral truths, which we don't derive from necessary moral truths of the type favoured by moral philosophers. Lord notes that it's implausible to think that necessary moral truths play a prominent role in the average person's acquisition of moral knowledge (Lord 2017, 18). Milona notes that it's unlikely that we support all our contingent moral knowledge with necessary moral truths (Milona 2017, 16-18).

I've outlined the ways in which expertise in moral judgment are different from expertise in moral action, analysis, and reasoning. For ease of discussion, I'll henceforth refer to expertise in moral judgment as 'moral expertise', and experts in moral judgment as 'moral experts'. Readers can assume that the following discussion of 'moral experts' and 'moral expertise' pertains exclusively to moral judgment, unless stated otherwise. The task is now to give a definition of both. Driver, no doubt, had an important qualification in mind when defining the moral expert in judgment. This is worth noting explicitly: the moral expert's correct moral judgments must be made in virtue of *her own* moral sensitivities. In other words, her expertise must derive from her own moral skill. We expect this much of experts in non-moral domains.⁴ So, for instance, someone who constantly relies on the moral testimony of others to arrive at correct moral judgments doesn't count as a moral expert.

Our definition of moral expertise, then, must stipulate that moral expertise

⁴ Howell notes that if one's surgeon comes into the operating theatre with a 'grin and a medical textbook', we wouldn't think much of him as a surgeon; at best, he's a practiced reader with a good pair of hands (Howell 2014, 400).

must exist in virtue of an agent's own moral sensitivities. This still isn't the whole picture, though. Since what we ultimately desire from experts is *knowledge*, it is worth building this explicitly into the definition of moral expertise. With all of this in mind, we can define moral expertise as the following:

Moral expertise: Moral expertise within moral sub-domain *D* is an agent's competence at forming moral knowledge in response to morally relevant features, within *D*.

2.3 Defining moral experts

Now that we've defined moral expertise, how can we define moral *experts*? This task is more complicated than it seems. We cannot simply define a moral expert as 'someone who has moral expertise'. First, note that moral expertise comes in *degrees*. It isn't a binary, on-off property of moral epistemic agents. One may have no moral expertise, a little of it, quite a bit of it, or an enormous amount of it. The fact that someone may have a low but real amount of moral expertise, then, shows that it's possible for one to have moral expertise without being a moral expert. There is a distinction between experts and expertise (Jones 1999, 64), and an agent's possession of moral expertise does not entail that they are a moral expert (Jones & Schroeter 2018, 461).

How then should we define moral experts? For now, I'll have to default to the

unsatisfying working definition that a moral expert is '*someone with a sufficiently high level of moral expertise*'. Unfortunately, this is also complicated by the fact that this sufficiency constraint can be read in two ways, depending on whether one holds experts to 'relative' or 'objective' standards. If standards for being an expert are *relative*, then whether or not someone is an expert in some domain depends on how competent the rest of the population is in that domain, or the competence of members of some contrast class.⁵ But if standards for being an expert are instead *objective*, then whether or not someone is an expert in some domain depends on whether they meet some population-independent standard of competence in that domain.⁶ I think that both notions of moral experts are useful. So I'll resist characterising my account solely in either set of terms. And on either definition of expertise, we can still surely identify clear cases of moral expertise and naivety at both extremes of the spectrum. So, my definition of a moral expert will suffice for our purposes.

⁵ One way to cash out relative expertise is to say that a moral expert is someone with more moral expertise than sufficiently high percentage of the general population. But we can also plausibly say that person (or group of people) A is a moral expert relative to person (or group of people) B, irrespective of where either of these agents stand in the general population. For instance, there is a sense in which *all* normal adults are moral experts relative to the class of five-year-old children (Driver 2013, 283).

⁶ If experts are held to an objective standard, then one (but not the only) possibility in is that moral experts are those whose rates of accurate moral judgment are sufficient to meet some pre-determined accuracy rate. A problem with this characterisation, though, is that it's difficult to stipulate what the precise cut-off point should be for expert status, or if there is any determinate cut-off point at all. Still, we can surely identify clear cases of moral expertise and naivety at both extremes of the spectrum. And we can surely say that one necessary requirement for someone to count as an objective moral expert is that their true moral judgments must count as knowledge.

3 The non-intellectualist picture of moral expertise

Relatively recent work on moral expertise by Driver, Jones, Schroeter, and others paint a combined, plausible picture of moral expertise, which stands in opposition with the ‘intellectualist’ picture painted by Singer. In this section, I’ll discuss this ‘non-intellectualist’ picture. I’ll begin by introducing three features of the current non-intellectualist picture (3.1). Next, I’ll propose a new addition to the non-intellectualist picture which has important implications for the rest of the paper. This is the claim that moral experts must be highly skilled at making contingent moral judgments (3.2). Last, I’ll note that the current picture lacks a good account of its most important claim—the claim that moral expertise is typically generated through experience (3.3).

3.1 Three features of the current non-intellectualist picture

The current non-intellectualist picture has argued for at least three features of moral expertise. I’ll outline them in turn.

(1) First, moral expertise is typically generated through concrete experience with the real world. Jones (1999) notes that rich experiences with certain types of moral problems can contribute to moral expertise. One way to have such richer experience is through one’s social location. Being a woman, for instance, can contribute greatly to one’s moral expertise on matters of sexism (65). One can also

have richer experiences in virtue of choice. For instance, someone who chooses to engage in political liberation movements is in a good position to gain moral expertise about social justice in those contexts (65-66). Driver (2006) agrees that experience matters, and notes that experience of the 'right sort' provides a plausible path towards moral expertise. Driver cites Mill's argument for the promotion of higher over lower pleasures to argue that sufficiently 'broad' experiences can be valuable. According to Mill, the only competent judges of whether 'higher pleasures' are superior to lower ones are those who have experienced *both* (627; Mill 1998, 56). Analogously, we have reason to give greater weight to the view of someone who has experienced both freedom and repression regarding which is to be morally promoted (628). In addition to the broadness of experience, Driver notes that we may also weigh the 'deepness' of experience. First personal experience, for instance, is often privileged. Someone who's experienced a tragedy is frequently taken to be a greater authority on that tragedy than someone who hasn't, since imagination, empathy, and sympathy only take you so far (628). Although experience is neither sufficient nor strictly necessary for moral expertise, it's certainly important, and the best way of thinking of experience is as providing a plausible mechanism for the development of moral sensitivity (629). Sreenivasan (2015), as a side note to his discussion on moral deference, notes that differences in experience may be a causal factor in moral learning; the 'wise people who populate legend and folklore are

invariably wise *old* people (46). McGrath (2011c) argues that being ‘acquainted’⁷ with some controversial practice often puts one in a better position to evaluate its morality, and briefly speculates that this puts pressure on us to incorporate conditions requiring acquaintance into an account of moral expertise. This stands in direct contrast with Singer’s intellectualist model of moral expertise, under which one could qualify as a full-fledged moral expert even in the complete absence of any acquaintance with those practices whose morality one is an expert (28).

(2) Second, moral expertise typically exists in local moral sub-domains,⁸ and does not exist globally over morality in general. Jones (1999) was the first to note the distinction between global and local claims to moral expertise, and argue that moral expertise is typically unevenly distributed with respect to certain domains (64-67). Driver (2006) also notes that it seems ‘far more likely’ that moral experts possess

⁷ For McGrath, one may be ‘acquainted’ with some real life event by witnessing it first hand, or seeing the event depicted in media like videos or photographs. In her sense of acquaintance, ‘one can be acquainted with a practice in virtue of viewing veridical pictures or videos of genuine instances of that practice. What distinguishes acquaintance with a practice from descriptive knowledge of that practice is that the former essentially requires a relation to a particular, concrete instance of the practice (even if that relationship is mediated in some way), while the latter does not’ (4, fn. 6). Like McGrath, I’m happy to count acquaintance with veridical pictures or videos as an important type of experience with the real world. Lord also invokes the term ‘acquaintance’ while discussing the role of experience, although he focuses on acquaintance with morally relevant properties rather than events, and thinks that one can also be acquainted with the moral features of non-veridical imaginative experiences (Lord 2017, 18-19).

⁸ Non-intellectualists haven’t explicitly provided any definition of what a ‘moral sub-domain’ is. The notion, while intuitive, is quite hard to pin down. I propose that moral facts can belong in the same moral sub-domain in virtue of at least three plausible reasons: (1) Agents being likely to encounter them in the same contexts, (2) invoking the same thick ethical concepts, and (3) being about the same purely descriptive phenomena. Notice that these three intuitive ways of identifying moral sub-domains support the idea that the way they’re individuated has a lot to do with how we *experience* the world.

'local expertise rather than global expertise' (625). Moral experts, far from always being geniuses who can discern rightness and wrongness in any type of situation, typically demonstrate expertise in specific moral sub-domains without necessarily demonstrating expertise in others. For instance, one may be an expert in detecting violations of free speech norms while failing to be an expert on sexism (625). This follows quite naturally from the fact that moral expertise is typically generated through experience, since different agents lead different lives and have varying experiences with the world. Jones & Schroeter (2012) note that since moral sub-domains 'present their own characteristic patterns in the combinations of moral considerations that tend to occur within them, ... [and] given that the world of value is complex and the capacities needed to navigate it are many and various, it is much more realistic to expect that such human moral expertise as it exists would take patchwork form rather than [an] idealised form' (223).

(3) Third, Jones & Schroeter (2012) have argued that moral experts do not always produce action guiding moral judgments. A good picture of moral expertise 'abandons the thought that expertise is shown only or primarily in all-in judgments about rightness and instead looks for more piece-meal context dependent expertise with particular thick moral concepts, such as *honest* and *respectful*, rather than with thin ethical concepts such as *right* or *ought* (219). All-in moral judgments 'form just a fraction of our moral judgments', and it is important for a good model of moral expertise to accommodate this (222). Once again, this follows naturally from the fact

that moral expertise is typically generated through experience. Since agents have different experiences, moral expertise typically exists in local sub-domains; this in turn means that moral expertise within some sub-domain does not guarantee expertise in making final action-guiding moral judgments about issues that cut across multiple sub-domains. What it does guarantee is expertise with the specific values present in local sub-domains.

3.2 The importance of contingent moral knowledge for moral experts

Now, I'll argue for a *fourth* feature that should be added to the non-intellectualist model, which has big implications for the rest of this paper. I argue that moral experts must be skilled at making *contingent* moral judgments about the actual world, and that their skill at making *necessary* moral judgments is much less important.

Necessary moral truths are moral propositions which are true in all metaphysically possible worlds, and necessary moral propositions are those which purport to be true across all metaphysically possible worlds. Candidates for necessary moral truths include 'murder is wrong' (Lord 2017, 18) and 'we ought to act so as to maximise (net) pleasure' (Milona MS, 4). If these claims turned out to be false, they would merely be necessary moral propositions.

Contingent moral truths, which I've briefly introduced in Section 1.2, are moral truths which are *not* true across all metaphysically possible worlds, and contingent

moral propositions are those which do not purport to be true across all metaphysically possible worlds. Schroeder gives an example of a contingent moral truth: 'It would be wrong for me not to show up for the tenth annual Wisconsin Metaethics Workshop'. This is contingent 'because I promised to attend, and hence could easily have been false, if only I had been selective about the commitments that I took on' (Schroeder 2014, 129). The truth of Schroeder's claim, of course, is contingent on more factors than his promise. For instance, a sudden family emergency might have dissolved his promissory obligation to attend the workshop. Or, the workshop could have been cancelled altogether. So the truth of contingent moral propositions hinges on multiple features of what the actual world is like, and the multiple values it presents.⁹ It's also important to note, as I have in Section 1.2, that we frequently gain knowledge of contingent moral truths *without* deriving them from necessary moral truths. Again, it's unlikely that we support all our contingent moral knowledge with necessary moral truths (Milona 2017, 16-18). And familiarity with moral theory and higher-order necessary principles doesn't seem necessary for moral knowledge (Tropman 2018, 474).

Any good model of moral experts, I argue, should prioritise their ability to gain contingent moral knowledge over necessary moral knowledge. In other words,

⁹ Another example is given by Milona, of the following contingent moral proposition: 'Bert was right to help the elderly man cross the street'. This proposition is contingent because, for instance, 'the elderly man could have been planning to poison his neighbour's friendly but noisy dog, in which case it wouldn't have been right for Bert to help (Milona 2017, 4). This proposition would be a contingent moral truth if the world exists in such a way that Bert is indeed right to help the elderly man (which depends on many more factors than whether the elderly man had plans to poison the dog).

moral experts should be good at gaining moral knowledge that is tethered to our everyday actual lives, and it is less important that they are good at acquiring necessary moral truths. I expect that many moral philosophers will resist this idea, since it stands in direct contrast with the field's predominant theoretical priorities. Traditionally, necessary moral truths are what ethicists quest after (Milona MS, 4). Investigating necessary moral truths has monopolised most theoretical discussions in the field, and the pursuit of contingent ethical truths is unpopular amongst moral philosophers (Lord 2017, 18).

But as Lord notes, how we think of moral epistemology shouldn't be shaped by moral philosophers' theoretical priorities (Lord 2017, 18). So, we shouldn't allow these priorities to shape what we want from moral expertise. First, as a matter of fact which I've noted in Section 1.2, a significant amount of moral knowledge (if not most of it) is of contingent moral truths. In addition to Lord (2017) and Milona (2017, MS), Schroeder (2014) has also noted that many of our important moral truths are contingent, rather than necessary (129). Second, this fact is not some undesirable symptom of our incapacity to easily acquire the necessary moral truths we're 'really looking for'. On the contrary, it accurately represents most of our moral interests in the world. To emphasise part of the point that Schroeder makes: many of our *important* moral truths are contingent, rather than necessary. Outside the field of moral philosophy, our everyday moral interests are about the moral status of things given the way things *are*. This partly explains why many students get frustrated

when their ethics professors present them with highly simplified, futuristic, hypothetical thought experiments. In some sense, their frustration is not misplaced. What we frequently want is to become good at assessing the complicated moral dilemmas, with numerous competing values, that are presented to us in the actual world. So, a good and useful account of moral experts must accommodate our everyday interests by requiring them to be skilled in making correct contingent moral judgments about the actual world, and the complicated moral cases it presents. If someone is not good at acquiring such contingent moral knowledge, then we cannot properly characterise her as a moral expert.

Non-intellectualists about moral expertise should be happy with this claim. First, the claim that moral experts should be good at acquiring contingent moral truths about the actual world rests well with the notion that our notion of moral expertise, and what we demand from moral experts, should be grounded in the real world and all its complexities. Another reason lies within their claim that moral experts will, and should, frequently produce expert moral judgments that aren't action-guiding, and instead take the form of whether thick ethical concepts apply. For instance, it is important that a moral expert in the sub-domain of racism should be able to identify something as *racist*. And whether or not thick ethical concepts like *racist* apply depends on contingent facts about the actual world. So, all such judgments are contingent moral judgments. If non-intellectualists wish to keep this feature of their model of moral expertise, they should be happy with the claim that a

moral expert should be skilled at making contingent moral judgments about the world. Last, as I will lay out in Section 3, the experience that non-intellectualists prioritise is a very good means of acquiring contingent moral truths about the world, and developing expertise in acquiring contingent moral knowledge.

3.3 The need for non-intellectualists to justify the role of experience

The non-intellectualist model of moral expertise is appealing, and the discussion above gestures us in the right direction. But what's still lacking is a detailed argument and moral epistemic story narrating precisely how particular experiences with the world can generate moral expertise—particularly moral expertise in acquiring contingent moral knowledge. When we experience the world, precisely what is it that happens that allows us to become better at knowing contingent moral facts in virtue of our own moral sensitivities? If we were to trace the development of someone who started as a moral novice and subsequently acquired moral expertise through her experience, how would that story be told? Precisely how do an agent's moral beliefs, and her capacity to acquire moral knowledge, change as she encounters the world? The current literature lacks such a detailed account.

Thankfully, some philosophers have given useful accounts of how experience helps us to acquire moral *knowledge*. Although these accounts stop short at addressing the issue of moral expertise, they do pave the way for the type of account I'm looking for. McGrath (2011a) argues that experience can play at least four crucial

roles in our acquisition of moral knowledge. Experience can contribute to moral knowledge by playing an enabling role, a triggering role, a sensitising role, and by supplying evidence for non-moral propositions that subsequently inform one's moral thinking. Lord (2017) argues that experience both puts one in an especially good position to gain moral knowledge about those particular experiences, and also plays an important role in the development of general moral sensitivities. These arguments about moral knowledge are instructive, and I will build on parts of their accounts in the section to come.

4 The role of experience in producing moral experts

Now, I'll tell the story that's missing in the literature on moral expertise. Using a sentimentalist perceptualist framework, I'll argue that extensive experience puts one in a good position to develop moral expertise, and eventually become a moral expert. This is because experience is good at causally prompting moral perceptions in the form of affective experiences, causally prompting contingent moral judgments about the actual world, and causally prompting moral adjustment over time. I propose that this can give rise to a form of moral perceptual learning, and this is how experience gives rise to moral expertise. What arises from repeated experiences with certain types of phenomena is an increased level of moral sensitivity with respect to those phenomena. This gives rise to an increased level of moral expertise in the

relevant moral sub-domains, and puts them in a good position to become a moral expert.

4.1 Preliminary notes about experience, sentimentalist perceptualism, and contingencies

First, here are some important notes about the picture I'll paint.

(1) When I say that experience puts one in a good position to develop moral expertise and become a moral expert, I mean that experience provides a path towards expertise that the typical agent can take advantage of under relatively agreeable circumstances. The possibility of becoming a moral expert depends on factors other than experience, such as one's affective constitution, reflective capacities, moral environment, and so on. Although I lack the space to fully address all the factors that make up these 'relatively agreeable circumstances', and the extent to which each of these factors matter, the reader should note that I'm not claiming that merely having experience is sufficient to generate moral expertise. What I mean is that experience puts agents who, for instance, are decently responsive to moral features and are not brainwashed by immoral ideologies in good positions to develop moral expertise.

(2) I'm also not making the strong claim that experience is strictly necessary for moral expertise and becoming a moral expert. Like Driver, I want to allow for the theoretical possibility that one may have moral expertise despite having

impoverished levels of experience; it is theoretically possible that there are five-year-old moral savants (Driver 2013, 283). What I will be arguing in Section 4, however, is that for the vast majority of agents in the real world, experience *will* turn out to be necessary after all. So if such moral savants exist at all, they are few and far between.

(3) Given the definition of moral expertise, this will be an account of just how experience puts one in a good position to acquire the moral sensitivities necessary for expertise. Here is the working definition of ‘moral sensitivities’ I will assume. An agent has expert moral sensitivities if and only if she responds to moral cases in some way that allows her to produce reliably correct moral judgments on her own. My definition of moral sensitivities and hence moral expertise, then, invokes reliabilism. However, I am open to the possibility that moral expertise is better captured in other terms, like that of virtue epistemology. The reader should not see me as being necessarily committed to reliabilism.

(4) My account will invoke, but not defend for reasons of space, the view of *moral perceptualism*. According to perceptualist accounts of moral epistemology, at least some moral properties can be part of the contents of perceptual experience. In other words, we can have ‘evaluative experiences’—or ‘evaluative intuitions’—which are a subset of perceptual experiences. Moreover, on the view of moral perceptualism I invoke, these evaluative experiences justify moral beliefs for the same reasons that ordinary perceptual experiences justify ordinary perceptual beliefs. Just as the contents of ordinary perceptual experiences can provide prima

facie justification for ordinary perceptual beliefs, moral properties present in the contents of evaluative experiences can provide prima facie justification for moral beliefs. My visual experience of seeing a red mug, for instance, provides prima facie basic justification for believing there's a red mug in front of me. In the same way, experiencing the moral intuition that one ought to save a drowning child from a shallow pond provides prima facie basic justification for believing that one indeed ought to do so.

(5) The variety of moral perceptualism I'll be invoking, but again not defending, is *sentimentalist perceptualism*.¹⁰ On this view, our moral perceptions come in the form of *affective experiences*, such as emotions and desires. It is affective experiences which form the foundations of our moral epistemology, and give us prima facie basic justification for our moral beliefs.

This stands in contrast with 'intellectualist perceptualism', which instead asserts that our moral perceptions come in the form of intellectual evaluative experiences (like the intellectual 'seemings' we get when we grasp mathematical truths). Sentimentalist perceptualism also stands in contrast with the moral perceptualist view of what's been called 'perceptual intuitionism'. On this view, moral properties can be part of the contents of experience in one of the traditional sense modalities like vision and hearing, as opposed to being part of the contents of

¹⁰ This type of view has been given various other names in the literature. For instance, it has been called 'response intuitionism' (Tropman 2018) and 'affectual intuitionism' (Roeser, 2011; Werner 2016). It has been defended by Döring (2007), Roeser (2011), and Milona (2016), discussed and not rejected by Cowan (2013; 2015) and Werner (2016), and invoked by Lord (2017) and Milona (2017; MS).

emotional experience (Werner 2016, 297).¹¹

Just how do affective experiences allow us to perceive moral properties? Sentimentalist perceptualists commonly note that many of our affective experiences involve experiences of value. For instance, anger involves a presentation of having been wronged, and desire involves a presentation of its object as good (Milona 2017, 4, 13). Another example from Döring is the indignation one might feel at the harsh punishment of a toddler who has accidentally dropped his ice cream. It seems to you that the punishment is unjust, and your affective state of indignation puts forth its content as correct. This is analogous to the content of a sense perception (Döring 2007, 377).

(6) I will follow sentimentalist perceptualists like Milona (2016; 2017; MS) and Lord (2017) in assuming that our affective experiences typically present *contingent* moral propositions as being true, as opposed to necessary moral propositions. And Lord notes, it is unpopular to think that we can learn about necessary moral truths via experience (Lord 2017, 18).¹² Milona argues that well-functioning affective experiences do not typically respond to contents which guarantee the presence of a corresponding value, and so the propositions that they help us know will typically be contingent, and not necessary (Milona MS, 17; 2017, 13-23). Zagzebski argues that

¹¹ See e.g. McBrayer (2010a; 2010b) for a defense of perceptual intuitionism. Werner (2016) notes that one reading of Audi (2013) can be interpreted as defending this view.

¹² More properly, Lord thinks that it is unpopular to think that we can learn about these truths via ‘acquaintance’. On Lord’s view, one can be acquainted with normative properties via both experience with the real world, as well as imaginative experiences. I agree with Lord, and discuss the role and limitations of imaginative experiences in Section 4.

an emotion is a state of affectively perceiving its object as falling under a *thick* affective concept (Zagzebski 2003). If this is true, then moral perception via emotions will always be sensitive to contingent descriptive factors of the situation, and will hence present contingent moral propositions to be true. On these views, the indignation one feels at the toddler's harsh punishment makes it seem to you that *that particular instance* of a harsh punishment was unjust, and not, for instance, that *all* harsh punishments are always unjust. And as Milona notes, this doesn't mean that we can 'stumble' into knowledge of necessary moral truths, or use our contingent moral judgments to justify necessary moral judgments (Milona 2017, 6, 20-23).

(7) Importantly, nothing about sentimentalist perceptualism implies that affective experiences are the *only* means through which we obtain moral knowledge. This is obviously implausible, one reason being that affective experiences, just like ordinary perceptual experiences, often mislead. Döring notes that one need not endorse the contents of one's moral perception, just as one need not endorse the contents of one's ordinary sense perceptions although they have the appearance of truth (Döring 2007, 378-379). The beliefs we form from moral perceptions are subject to adjustment, and indeed, refining one's beliefs by learning how to use one's moral perceptions is the gateway to moral expertise and becoming a moral expert. This will be made clear in the discussion to come.

With these qualifications out of the way, I'll now proceed to show how

experience with the real world puts one in a good position to develop moral expertise.

4.2 Experience causes moral perceptions and immediate moral judgments

First and foremost, experience plays the role of causally triggering affective responses to normative properties. Taking an example from Lord, someone who witnesses a hanging for the first time will probably experience a host of affective reactions to the scene in front of him, such as disgust, repulsion, and sadness (Lord 2017, 10). Such affective experiences come prior to moral judgments, and agents may have them without coming to any moral judgments at all.¹³ For one, as previously noted, one need not endorse one's immediate moral perceptions as being true. Moreover, an agent may not recognise the moral implications of her affective experience. For instance, in a series of real-life interviews of subjects who have witnessed prisoners being executed via lethal injection, many interviewees reported a host of negative affective reactions to the experience of watching someone being killed. None, however, reported making the explicit judgment that the executions they witnessed were morally wrong.¹⁴ I take it that it is extremely common for

¹³ Here, I'm assuming the view that affective responses that form the bases of our beliefs are prior to those moral beliefs (see, e.g. Audi 2013). But there are other views which claim that affective responses aren't prior to moral beliefs, but instead combine with them to form a single unitary state. Roser (2011), for instance, argues that moral emotions are candidates for genuine moral knowledge in this way.

¹⁴ From an interview by the New York Times, 'Bearing Witness to Executions: Last Breaths and Lasting Impressions.' (April 23, 2017)

agents to have affective moral perceptions without yet endorsing any moral judgments that these perceptions might justify, or even recognising that they are engaging in anything like moral perception at all. Experience, then, is a good way to trigger moral perceptions in the form of affective reactions, so long as moral agents are decently affectively responsive to at least some moral features of the world.

After experience prompts an agent to have an affective experience that constitutes a moral perception, her next moral epistemic step is to use that perception to form a moral judgment (or set of moral judgments). One way for her to do this is to take her affective experiences at face value, and immediately endorse the contents of her moral perceptions. For such an agent, experience causally contributes to the formation of moral judgments quite directly through her moral perception. Lord notes in his hanging example that the host of negative affective reactions one may have while witnessing a hanging can allow one to learn and judge that the particular hanging is wrong, and that hanging in general is wrong (Lord 2017, 10). Döring notes that an emotion can lead an agent to judge a content involving an evaluative notion, that he would not otherwise in fact have judged (Döring 2007, 390). One simple way to read both Lord and Döring's examples is that such agents are taking their affective experiences at face value. McGrath notes that experience plays a psychological role in prompting one to make moral judgments one would otherwise not have made (2011b, 7). She remains silent on just how one may come to form this judgment; on my account, this might happen when an agent takes her

affective moral perceptions at face value.¹⁵ Experience is a good way to trigger these sorts of moral judgments, since experience is an especially good way to trigger the moral perceptions that form the direct basis for these judgments. Importantly, note that the moral judgments that moral perceptions directly provoke will typically be *contingent* judgments about the the particular situations that provoked those perceptions.

4.3 Experience causally prompts moral adjustment

An agent, of course, need not take her moral perceptions at face value. Although most of us do this from time to time, I take it that only unwise or morally naïve agents are in the habit of *always* taking their affect at moral face value. So, what I've said so far is a long way from establishing the link between experience and moral expertise. Here, I'll take us one step further by arguing that experience puts one in a good position to undergo what I call moral adjustment.

Instead of taking their moral perceptions at face value, agents may instead compare their occurrent perceptions against prior moral beliefs, or other conflicting occurrent moral perceptions. Audi notes that following an affective moral intuition, we may choose to withhold judgment and engage in 'reflection', which can confirm or disconfirm our initial intuition, providing a path to an eventual moral judgment

¹⁵ I've simplified things somewhat in my citations of Lord, Döring, and McGrath. The most straightforward way to read them (or at least the sections I've referenced) is to see them as addressing how agents can unreflectively make moral judgments on the spot. But their accounts can certainly accommodate intermediary reflection between having a moral perception and making a moral judgment.

(Audi 2013, 165). Unlike Audi, I'll resist using the word 'reflection' to exhaustively characterise this type of process by which we can arrive at more sophisticated moral judgments. This is because it risks over-intellectualising the possible ways in which one might refine one's moral beliefs. I wish to allow for the plausible possibility that this can happen without too much explicit deliberation on the agent's part. Consequently, I'll refer to this process as *moral adjustment*. The process of moral adjustment includes, but isn't exhausted by, rigorous and conscious moral reflection.

For the sake of illustration, consider an agent, Kelly, who has so far only ever taken her moral perceptions at face value. Due to her particular experiences and perceptions, the moral judgments she's formed and the moral perceptions she's had have never come into conflict. Although Kelly seems fortunate in some respect, the consistency of her moral perceptions and judgments doesn't guarantee their truth. We think she would do well with some extra adjustment to arrive at more sophisticated moral judgments. How might experience help Kelly do this?

First, note how unrealistic an agent like Kelly seems. Our experience of the world is complicated, and few of us have never had reason to question our prior moral beliefs, or immediate moral intuitions. The more we experience the world, the more we encounter complex situations involving multiple values, where either our moral perceptions conflict with our prior moral judgments, or our occurrent moral perceptions conflict with each other. In other words, experience tends to present us with complex moral dilemmas. The need to resolve this uncertainty is often what

prompts us to engage in further moral adjustment.

Let's return to naïve Kelly. Until today, she has only encountered cases of stealing which are straightforwardly morally wrong. Once, she had her purse snatched away by a stranger on the street. When this happened, Kelly took her indignation and anger to form and justify a belief that she had been wronged by the thief when he stole her purse. Some time after, she witnessed a well-dressed teenager slip a candy bar into his pocket at a small family-owned store. Again feeling indignant, Kelly took this emotion to justify a belief that this instance of stealing was wrong. Moreover, she may have taken both of these experiences as justification for the necessary moral proposition that 'stealing is always wrong'.

Today, however, Kelly witnesses a man in tattered clothes being arrested at a grocery store for stealing infant formula. As he's being dragged away, the man protests loudly that he had no choice but to commit the crime. As a single father who's just been let go from his job, he's had no money to feed his baby. Kelly finds the man's testimony compelling. Far from experiencing indignation at the fact that infant formula has been stolen, Kelly now feels sad at the man's predicament, and outrage that he is being arrested for attempting to feed his child. She believes this is something he is morally compelled to do, and that others are morally compelled to allow him to do so. Kelly comes to the contingent moral judgment that this instance of stealing was morally permissible, if not morally required. (She will also plausibly realise that this conflicts with her prior necessary moral judgment that 'stealing is

always wrong', and updates her belief in light of this, to form new necessary moral judgments like 'stealing is sometimes permissible', and 'it is morally permissible to steal infant formula if a child's life is at stake'.)

In the case above, experience has prompted Kelly to arrive at more nuanced contingent (and necessary) moral judgments about the permissibility and wrongness of stealing in the face of competing value. Experience, however, might also prompt Kelly to refine the *scope* of cases she takes to count as stealing at all. Assuming that Kelly takes 'stealing' to be a thick ethical concept, experience may prompt her to refine the scope of cases she's willing to apply this concept to. Imagine that Kelly now sees on the news that someone has been caught taking valuable artefacts from her favourite British museum in the middle of the night. Kelly feels indignation at the man's actions, and pleasure at his being caught; she consequently judges that justice has been done to someone who has done something wrong by stealing valuable property. However, the news anchor proceeds to note that this man is an activist, who took the artefacts with the intention with returning them to Indonesia. These artefacts were taken from an Indonesian family by Dutch colonisers in the 1820s, without payment or consent, and eventually sold for great profit to the British museum. Moreover, it turns out that the activist is a descendant of the original family from whom the artefact was stolen. Kelly, who happens to have strong feelings about colonialism, now feels outraged that the activist is being punished for taking back something which was stolen from his ancestors. One thing she might do

is to use her affect to conclude that this is a morally permissible case of stealing. However, another alternative—which Kelly ends up taking—is to conclude that what he did was *not* stealing, since the artefacts in question had originally been stolen from his ancestors, and what he did was simply to take back what was rightfully his in virtue of his heritage. Experiencing this complicated moral case, then, has prompted Kelly to refine the scope of what she's willing to apply the thick ethical concept of 'stealing' to. Moreover, it has plausibly prompted her to refine how's she's willing to apply thick ethical concepts related to stealing, such as 'property', and prompted contingent moral judgments about the property rights of the various parties involved in the case.

The cases above demonstrate that experience can cause one to have moral perceptions that conflict with prior moral beliefs, and hence easily prompt moral adjustment to produce more sophisticated moral judgments. But note that it's also possible for moral adjustment to be prompted by experiences that cause one to have conflicting occurrent moral perceptions about the same case. Indeed, this does seem like part of what's happening to Kelly in the museum case above. Upon hearing the news, Kelly has a host of conflicting affective experiences in a matter of seconds: initial indignation at the thief, pleasure at his being caught, subsequent indignation at the actions of the Dutch colonisers, admiration for the activist, and so on. Not only does she experience psychological tension between some of her moral perceptions and her formerly-held beliefs about stealing, she also experiences psychological

conflict between her moral perceptions in the moment.

To sum up, what the discussion above shows is that experience puts agents in a good position to engage in moral adjustment in order to arrive at more sophisticated contingent moral judgments (as well as to obtain some amount of justification for certain necessary moral judgments). This is because experience tends to present us with complex cases which produce moral uncertainty, which is a good means of prompting moral adjustment.

4.4 Experience giving rise to moral perceptual learning

Let's assume that everything Kelly's concluded so far is correct. What has experience handed her? One thing she's obviously gained is new moral knowledge. She now knows that 'stealing can sometimes be permissible', and 'what the activist did wasn't an instance of stealing'. But more importantly for our purposes, Kelly's been put in a good position to gain moral *expertise*, and to eventually become a moral expert with more experience. Experience has put her in a good position to improve her *ability* to acquire moral knowledge in virtue of her own sensitivities. Moreover, she's in a good position to acquire *contingent* moral knowledge, and is thus in a good position to become as the type of moral expert we want.

I propose that Kelly is in a good position to have undergone, or to undergo with repeated experience, the moral analogue of what philosophers have called

perceptual learning.¹⁶ Perceptual learning refers to ‘long-lasting changes in perception that result from practice or experience’, and such cases have often been used to show that ‘through learning we come to represent new properties in perception, which we did not represent prior to learning’ (Connolly 2017). Much of the philosophical discussion on perceptual learning has focused on vision. Siegel, for instance, argues that spending months cutting down only pine trees in a particular grove of diverse trees can cause pine trees to look different to us. At the start of this exercise, the pine-tree identifier must infer whether or not some tree is a pine tree from various physical traits of the tree. At the end of this exercise, Siegel argues, they can come to see pine trees *as pine trees*, and need not infer this from any more basic physical characteristics of the trees in question. This is a case of perceptual learning, where perception comes to represent the property of ‘being a pine tree’, whereas it did not represent this before. The property of ‘being a pine tree’ becomes part of the content of one’s perception (Siegel 2006; 2010). Some philosophers have also explicitly discussed perceptual learning in experts (see e.g. Siewert 1998; Raftopoulos & Zeimbekis 2015).¹⁷ Driver notes how an expert zoologist might see that a particular swirl in the sand has the property of a ‘being a rattle-snake mark’, while her daughter might simply see it as being a swirl in the sand. The zoologist might also

¹⁶ Thanks to Errol Lord for this suggestion.

¹⁷ Siewert notes that a chessboard in midgame might look differently to a chess player than a novice, and a car engine might look differently to a mechanic than to someone unfamiliar with cars (Siewert 1998). Raftopoulos & Zeimbekis consider a scientist who, through repeated exposure to items in her expert domain, has developed perceptual sensitivity to certain features. As a result, she might quite literally see the world differently within her expert domain than a non-expert (Raftopoulos & Zeimbekis 2015, 19).

see that a bird nesting in a particular cactus has the property of ‘being a rare species’, while her daughter simply sees it as a bird nesting in a cactus. According to Driver, the zoologist is ‘able to make finer discriminations due to her expert knowledge, and this in turn really does influence how she sees the world around her’ (Driver 2013, 295).

In all these cases, experience has allowed agents to develop their basic visual perceptual capacities to a higher level. In other words, experience has allowed these agents to gain new visual perceptual expertise. Someone who is capable of non-inferentially visually representing the property of ‘pineness’ has expertise in identifying pine trees, and someone who is capable non-inferentially of visually representing the property of ‘being a rattle-snake mark’ has expertise in identifying rattle-snake marks. Through repeated exposure to pine trees and rattle-snake marks, these agents have learned to use their basic visual capacities to non-inferentially represent artefacts in the contents of their visual experience.

For moral perceptualists, it’s plausible that something analogous to visual (or other types of) perceptual learning. In other words, we can undergo *moral perceptual learning*. If we have moral perceptual capacities, and if the rest of our ordinary perceptual capacities are subject to perceptual learning, then it’s natural to think that we can go moral perceptual learning as well. Although Goldie (on one interpretation) doesn’t subscribe to my type of sentimentalist perceptualism,¹⁸ he

¹⁸ For the most part, Goldie does not explicitly address what exactly he means when he says experience and emotion can train us to ‘literally see’ what ought to be done, or what thick

makes the analogy between ordinary perceptual learning (and skill) and moral perceptual learning (and skill), quite explicitly:

‘Some people, such as those with certain [non-moral] skills or [moral] virtues, can, because of their ability, perceive non-inferentially what others, with normal eyesight, hearing, and so on cannot. Thus two people can be looking in the same direction, at the same part of the scene, and yet one—the [non-morally] skilled or [morally] virtuous person—sees things that the other fails to see. And, once we have room for this, we can add the further point that [non-morally] skilled and [morally] virtuous people can be trained, developing an ability that can be learned ... This expertise will come in degrees.’ (Goldie 2007, 349)¹⁹

evaluative concepts apply to some case. I find it quite unclear as to precisely what moral epistemic view he espouses. But he briefly notes that the moral expert’s feeling of ‘shame’ at looking through someone’s diary is not a justification for the belief that the action would have been (wrongful) ‘prying’. ‘Rather, both feeling and belief are justified by further reasons, such as, perhaps, the fact that the action would have been demeaning to both parties.’ (Goldie 2007, 356). On one reading, this indicates that he thinks affective experiences cannot provide basic justification for moral beliefs.

¹⁹ What Goldie says about the ‘virtuous person’ happens to fit quite nicely with the non-intellectualist picture of moral expertise. According to Goldie, it is implausible that most of us could ever achieve the ‘unity of the virtues’, or Aristotle’s sense of ‘practical wisdom’ (*phronēsis*), where one has a ‘single complex sensitivity’ that allows one to perceive non-inferentially what the *right* thing is to do. Agents typically do not have ‘all’ the virtues, and excel in respect to only a few. So, a morally skilled agent with some virtue (or set of virtues) will only perceive the demands of that virtue (or that set of virtues) (Goldie 2007, 358-360). This fits quite nicely with the non-intellectualist’s claim that moral expertise typically exists in local moral sub-domains, and that expert moral judgments are often non-action-guiding.

For sentimentalist perceptualists like me, the perceptual capacity which undergoes learning will be our *affective* moral perceptual capacity. Through repeated exposure to various types of moral phenomena, agents can learn to use their basic affective moral capacities to represent more sophisticated normative features in the very contents of their affective experiences.

Let's go back to Kelly, and how she's been put in a good position to undergo moral perceptual learning. So far, she has used her basic moral affective capacities engage in some type of adjustment or inference that what the museum thief did was not stealing. But if Kelly continues to be repeatedly exposed to cases of apparent stealing—let's say she takes on an internship at her local police department—she's in a good position to undergo moral perceptual learning by refining her affective capacities with respect to stealing and property rights, such that she no longer has to engage in any moral inference or adjustment at all. It looks like Kelly already has the moral perceptual expertise to perceive clear-cut cases of stealing *as stealing*. But repeated exposure plausibly puts her in a good position to accurately perceive more complex cases *as stealing* or *not stealing*, unlike in the museum case where her conclusion that what the thief did was not stealing was inferential. Repeated exposure will also plausibly put her in a good position to accurately perceive cases *as a permissible case of stealing*, or *an instance of wrongful punishment for what is stealing*. She may also become good at perceiving what *the compassionate thing to do* is when a thief has been caught.

In Goldie's example, one gains the ability to literally perceive what the *kind* thing is to do²⁰ through repeated experiences which call for being kind (Goldie 2007). McGrath gives the example of someone who has participated in many close friendships; as a result of her experience she's in a good position to recognise which disclosures of a friend's personal information would constitute betrayals, and which would not (McGrath 2011a, 19). On my account, such an agent may come to be good at representing such cases *as betrayals* in her affective experiences. Through repeated exposure, agents are put in good positions to upgrade their moral perceptual capacities to accurately make these types of contingent moral judgments. Like the expert pine tree identifier, they can now become experts at detecting these contingent moral features by being capable of accurately representing them in her affective perceptions.

One might make the following objection to my proposal, which rests on a common objection to sentimentalist perceptualism in general. This is the observation that moral beliefs don't always appear to be accompanied by feelings. As Tropman notes, it seems possible for an agent to believe sincerely that something is morally bad without experiencing any kind of emotional response (Tropman 2018, 481). If expert moral judgments are the result of a fine-tuned affective system, how do we account for such judgments? My first response to this is to note that types of moral

²⁰ More properly, what's happening here is that the agent perceives 'what the kind thing is to do' by engaging in counterfactual imagination. This perception itself happens in response to the imaginative experience of the counterfactual scenario provoked by the agent's moral dilemma in the concrete world. I will address this in Section 4.3.

judgments we tend to make ‘dispassionately’ tend to be about *necessary* moral propositions, such as ‘one has a moral duty to keep one’s promises’ and ‘one has a duty to take others’ ends as one’s own’ (Tropman 2018, 481) or ‘one always ought to maximise (net) pleasure’. But, as I’ve argued, it is much more important for moral experts to be capable of making accurate *contingent* moral judgments, which are much more likely to provoke affect. Zagzebski claims that it is possible to judge that ‘something is pitiful without feeling pity’ judge that ‘something is rude without feeling offended’ (Zagzebski 2003, 120). Yet, to me at least, it really seems quite implausible to *sincerely* make these moral judgments without the accompanying affect. Zagzebski notes the importance of being able to speak about concepts like ‘rudeness’ and ‘pity’ in a dispassionate way. I agree that it is good for us to be able to make claims like ‘something is rude if...’ and ‘someone is pitiful if...’. These are the best candidates for ways in which we may talk about such concepts in a dispassionate way. Yet they are *necessary* moral judgments about the definitions of such concepts; they are not contingent moral judgments about particular instances of rudeness and pity which are likely to provoke affect.

Moreover, nothing I’ve said blocks off the possibility that processes other than affective moral perceptual learning contribute to moral expertise as well. There are, for instance, some very instinctive emotional reactions that are hard for us to unlearn, such as fear and disgust, and moral experts should know when to distrust them. Milona argues that agents would do well to learn to recognise when their

evaluative perceptions are misleading, and learn to wisely override them (Milona 2016, 902-905). It is possible, and plausible that learning how to do this in non-inferential ways is important for moral expertise, even though it may involve no change in the contents of one's affective perceptions.

I lack the space to fully defend the proposal that moral expertise can be gained via moral perceptual learning, and acknowledge that it is a controversial proposal. But I think it is an promising possibility that sentimentalists perceptualists should seriously investigate, and one that all moral perceptualist should investigate in similar forms, at least if they wish to preserve a tight analogy between moral epistemology and ordinary perception. If we have moral perceptual faculties which are analogous to our ordinary perceptual faculties (which moral perceptualists hold), and if our ordinary perceptual faculties are subject to perceptual learning (which is plausible), then it is likely that our moral perceptual faculties are subject to perceptual learning as well.

5 The challenges of imagination

So far, I've argued that having extensive experiences with the real world provides a good way for one to acquire moral expertise and become a moral expert. I still need to show how experience provides an *especially* good way for one to acquire moral expertise. That is, I need to show how it is typically superior to other potential

means of acquiring moral expertise that do not involve experiencing the real world.

It's now time to address what some will see as the elephant in the room: *imagination*. I take imagination to be the main contender against concrete experience of the real world as a source of moral expertise. This worry hasn't been addressed by non-intellectualists. But now that I've invoked sentimentalist perceptualism, it is especially pressing for me to address the issue. One reason is that the sentimentalist perceptualists whose views I've been relying on have argued that using our imagination to get contingent moral knowledge is at least as good a means as using real world experience. What they primarily have in mind are the *armchair thought experiments* that are rampant in the methodology of moral philosophy. Here is Lord:

'The most common way we acquire armchair knowledge of contingent truths is by thinking about *cases*. When we do this, we imagine the cases in at least some detail. The circumstances that we imagine have certain morally relevant properties. ... I think it's plausible that we often have affective perceptual experiences of these properties. Thus, I think that ultimately the way we acquire knowledge of contingent truths via thinking about cases is very similar to the way we acquire knowledge of contingent truths by encountering concrete situations in the actual world.' (Lord 2017, 18-19)

And here is Milona, on the parity of 'online' and 'offline' moral experiences:

‘The moral experiences that we have in response to our perceptions and beliefs about the world—our *online* moral experiences—do not seem to be privileged (or at least not always privileged) over the moral experiences that we have in response to suppositions and imaginings—our *offline* moral experiences. ... While we rely on *actual* experiments to figure out the nature of the empirical world, moral inquiry only seems to require *thought* experiments. ... Offline affect is incredibly similar to online affect, and the former plausibly gets its content by virtue of being a simple extension of the latter.’ (Milona MS, 3, 19)

Other philosophers who aren’t sentimentalist perceptualists have also noted the power of imagination to provoke morally important affective experiences (e.g. see Audi 2013;²¹ Goldie 2007).²² And beyond armchair thought experiments, many others have also noted that imaginative engagement with *narrative fiction* can inspire important moral emotions (see e.g. Nussbaum 1990; Goldie 2007; Johnson 2016; and

²¹ Audi (2013, 157-161): ‘Intuition can arise from imaginative experience and can indeed provide evidence for moral judgments in those cases as well as where emotion arises from actual perception. ... We can sometimes be as emotionally sensitive, or as intuitively insightful, when we imagine a situation calling for a decision as when we actually see one person relating to another in an actual situation of just that kind. ... The exercise of moral imagination can, through vivid imaging of morally significant events, and through envisaging diverse possibilities, produce an experience significantly like a moral perception. ... Moral imagination, then, can provide evidence and do so in a way that parallels the way moral perception provides it.’

²² Goldie (2007, 354-356) gives a discussion of how emotions provoked by counterfactual thinking can contribute to the development of virtue and moral skill.

others).²³ As Johnson notes:

‘The importance of moral imagination helps account for the large number of philosophers, literary theorists, and psychologists who argue that moral cultivation comes more properly from our engagement with historical and fictional narratives than from treatises of moral philosophy that seek to articulate either absolute moral principles, catalogues of the virtues, or ultimate conceptions of the good. It is the narrative depth, complexity, and the existential validity of literary fictions that situates our moral perception²⁴ and appraisal in contexts that are psychologically fictional. Narratives can be morally transformative, insofar as we come to imaginatively inhabit the world and lives of characters that are both like and unlike ourselves and the people we encounter in our lives.’ (Johnson 2016, 365)

The threat of imagination becomes even worse when we consider the fact that it gives us access to seemingly boundless possibilities. Concrete experiences are limited in a way that imagination is not. Audi notes:

‘By contrast with perception [in the actual world], imagination is creative ... it

²³ Also see Gardner (1978), Eldridge (1989), Gregory (2009), Friend (2016), and Kind (FC).

²⁴ It’s doubtful Johnson espouses the type of moral perceptualism invoked by Audi, Lord, Milona, and me. Here, he likely invokes a more metaphorical sense of ‘moral perception’ which is more akin to something like ‘moral sensitivity’.

is properly limited by (at most) possibilities. Imagination apparently depends on the world, or at least on experience, for raw material; but it can build indefinitely many structures from that material ... Perception [of the actual world] is tied to its objects; imagination, even if it requires raw material from perception, is limitlessly combinatory, often dynamic, and readily responsive to our desire for even minute alteration. In the moral realm, imagination can construct morally significant scenarios we have never experienced.' (Audi 2013, 158-159)

It would be ideal to get a clear grip on precisely *what* imagination is, or at least what type of imagination is being invoked by the authors above. Unfortunately, the literature on imagination is especially murky.²⁵ So, I'll begin by assuming a rough functional characterisation of the type of imagination important for moral development. Here are three plausible necessary conditions. If an agent is engaged in *morally important imagination*, then (1) the contents of her mental state have no immediate input of sensory experiences of the real world, (2) she is not dreaming or hallucinating, and (3) her mental state provokes or is capable of provoking moral perceptions in the form of affective responses. This rough characterisation will suffice for our purposes.

²⁵ There is a great deal of disagreement amongst philosophers and psychologists about what imagination is, its metaphysical status, what mental states can properly be counted as imaginative, its phenomenology, and so on. Moreover, Gendler (2016) notes that the general consensus amongst those in the field is that the term 'imagination' is used too broadly to permit simple taxonomy.

5.1 The challenges of thought experiments and narrative fiction

Both the 'narrative fiction' and 'thought experiment' camps I've described observe the importance of imagination in moral epistemology. But interestingly, they threaten to stand in tension with each other. Traditional thought experiments are liberal about what sorts of details and particularities may be abstracted away, or how unrealistic the cases described may be. The whole point of such thought experiments, it is commonly thought, is to direct one's attention to morally salient features by cutting out irrelevant and distracting details. Moreover, philosophers frequently take advantage of imagination's boundless possibilities to do so, and thought experiments often describe wildly implausible scenarios.

Yet for those who stress the importance of fictional narratives in moral development, fiction is supposed to do quite the opposite. Nussbaum explicitly notes this opposition by saying that philosophers' examples almost always lack the particularity, the emotive appeal, the absorbing plottedness, the variety and indeterminacy, of good fiction' (Nussbaum 1990, 46). So, good fiction comes as close to approximating real life as possible. And the moral power of such fiction comes from getting us as close to *experiencing* these realistic narratives as possible. Note, from the quote above, that Johnson thinks fiction can morally transform us insofar as we come to 'imaginatively inhabit the world and lives of characters'. And 'moral imagination, both as our capacity to empathically understand and feel with others

and our ability to imagine how experience would play out under the shaping influence of various values and choices, is thus dependent on our ability to *simulate experiences* ... [allowing] us to experience something of the feelings and emotional responses of other people' (Johnson 2016, 364, italics by me).

Philosophers like Milona and Lord, of course, don't disagree that the moral power of imagination derives from its ability to 'simulate experiences'. Milona speaks of 'sensory imaginings' (Milona MS, 16) which provoke moral perceptions, and Lord uses the term 'imaginative experiences' (Lord 2017, 18). This suggests that morally powerful imagination is simply a type of non-veridical experience. Where these two camps come apart is on how close to concrete experiences these simulations should be. Those in the thought experiment camp tend to be liberal about cutting out what they see as irrelevant particularities, and invoking unrealistic but morally relevant detail, when engaging in moral imagination. In this sense our imagination is permitted to be far removed from concrete experience. Those in the fictional narrative camp think that the moral power of imagination derives from its close approximation to reality, along with all the particularities and indeterminacies of concrete experience.

It's also important to note one may think that *both* thought experiments and fictional narratives are morally powerful. The ways in which they are powerful are simply different. Lord, for example, has had undergraduates read both works of literary fiction and thought experiments in introductory ethics classes (Lord 2016). A

given work may also simultaneously invoke the powers of both. Good thought experiments often derive some of their moral power from carefully chosen particular details. And fictional narratives may also derive some of their moral power by situating their characters and plots in futuristic settings, or be morally powerful in spite of their fantastical nature. Science fiction, for instance, can be morally powerful. There are also fictional works that bluntly highlight moral features in the way that thought experiments tend to do. Episodes of *Star Trek*, for instance, are often (artistically) guilty of presenting crude plots that hammer in moral dilemmas. And writers have told moral tales that are difficult to properly classify as either a thought experiment or traditional literature (see e.g. Le Guin 1973).

For my purposes, however, I will stick to the simpler dichotomy that I've set up. Both traditional thought experiments and traditional fiction present challenges to the privileged role of concrete experience in the development of moral expertise. I will tackle each of them in turn. What I have to say about each of these challenges can also be applied to 'in-between' cases. First, I'll tackle the challenges of thought experiments, considering both the philosopher's method of deliberately engaging in lots of time conducting them (5.2), and the method of engaging in thought experiments under the constraints of ordinary life (5.3). Next, I'll tackle the challenges of narrative fiction, first considering the method of engaging in fiction under the constraints of ordinary life (5.4), and also the method of deliberately spending lots of time immersing oneself in fiction (5.5).

5.2 The challenge of thought experiments: round 1

As it turns out, I agree that thought experiments are a useful and powerful tool in our moral epistemic kit. One cannot deny their force; many college students would undoubtedly never have gained the basic piece of moral knowledge that ‘one is morally required to save a drowning child from a pond at relatively little cost to oneself’ if not for Singer’s famous thought experiment (Singer 1972b). It is also certainly good that thought experiments allow us to construct hypothetical cases that one would not encounter otherwise. And I agree with Lord and Milona that imaginative thought experiments, used properly, can elicit moral perceptions that are identical in type with those elicited by real life experience.

What I’m skeptical of is the claim that merely engaging in thought experiments is typically a good means of developing *high levels of moral expertise* in specific moral sub-domains, that allow us to make nuanced contingent moral judgments in relatively complicated solutions. Imagination alone cannot replace the years of targeted concrete experience that gives rise to this level of expertise. There are at least two ways of comparing engaging in thought experiments against having concrete experiences. The first is to evaluate the method of deliberately spending lots of time doing so (i.e. the moral philosopher’s method). The second is to evaluate whether for the vast majority of the population who are not professional ethicists, the moral epistemic value of how much time and effort one can reasonably spend

engaging in thought experiments can compare to that of the time one spends effortlessly experiencing the real world. I will begin by evaluating the first.

One question we can ask is what limitations our concrete experience place on our moral imaginative capacities, no matter how much time we spend engaging in thought experiments. Audi acknowledges that imagination ‘depends on the world, or at least on experience, for raw material’ (Audi 2013, 158). On one basic interpretation, this is just to agree with Descartes’ famous comment that ‘even when painters try to depict sirens and satyrs with the most extraordinary bodies, they simply jumble up the limbs of different kinds of real animals, rather than inventing natures that are entirely new’ (Descartes 1641, 2). In other words, our sensory experiences with the real world limit what we can depict in non-veridical mental states like imagining and dreaming.

I grant, however, that the concrete sensory experiences of even the most hermit-like ethicists are sufficient raw material for constructing decent moral stories. Moreover, we do not think that blindness or deafness poses any limits on someone’s ability to be a good moral philosopher. Is the capacity to engage in thought experiments limited by concrete experience in any other way? One more realistic possibility is to say that imagination is limited by the higher-level types of events one has experienced, or otherwise learned about, in the real world. In other words, the narratives one can concoct, and the moral dilemmas one can pose to oneself, are

limited by the phenomena one has witnessed or experienced in the real world.²⁶ For instance, Singer could not have independently imagined a child drowning in a shallow pond, or having to decide what to do about this, if he was unaware that people can drown. Philippa Foot could not have independently imagined having to choose between diverting a trolley to save five lives and kill one person, or letting it run its course to leave one person alive and let five die, if she was unaware that trolleys can kill, or unaware that trolleys can be diverted (Foot 1967). But this kind of limitation can be overcome if one has enough creativity to construct analogous cases of the same type, if one already has that type of moral case in mind.²⁷ Again, I am willing to grant that most moral philosophers are creative enough to do so. This is supported by the fantastical nature that many of their thought experiments have.

At some point, however, an agent's lack of life experiences can hinder her imagination in such a way that she is blinded to some types of moral scenarios altogether. One can imagine that Foot's concrete experiences might have been limited to such an extent that she could not have independently conceived of the moral distinction between killing and letting die. If so, no amount of narrative creativity could have allowed her to philosophically deploy a thought experiment

²⁶ This ability is somewhat related to what Van Leewuwen calls 'constructive imagination', which 'generates representations of possible states of the world—ways the world might be—on the basis of which we might choose actions', and 'generates representations of possible actions to take' (Van Leewuwen 2016, 295)

²⁷ For instance, Singer could have concocted some non-drowning story demonstrating the obligation to help at relatively little cost to ourselves. And Foot could have concocted some non-trolley story demonstrating the difference between killing and letting die, and having to consider the uneven consequences of some set of choices.

analogous to the trolley problem. And this brings me to the third and much more substantive limitation that lacking real life experience imposes on one's imaginative ability to engage in thought experiments. What we are able to conceive of as morally relevant features are influenced by our concrete experiences of the world. We do not construct thought experiments by churning out all possibilities via brute force, picking out the most interesting ones. We construct them on the basis of morally relevant features that we already have in mind.²⁸ Speculative as this may be, perhaps it is no accident that Judith Jarvis Thomson—a woman—was the first participant in the previously male-dominated philosophical debate on abortion to even identify women's bodies as morally relevant to abortion at all, and demonstrate this glaring point by constructing her famous violinist thought experiment (Thomson 1971). As creative as using 'the possibility of transforming kittens into persons' to demonstrate a foetus's non-right to life,²⁹ such ingenuity cannot make up for the ability to identify morally relevant features which experience makes obvious, but are nonetheless obscured from one by lack of it.

It may turn out, of course, that Thomson being a woman had nothing to do with any of this. But there are other examples. Consider the moral issue of late-term abortions. As a matter of fact, many of these pregnancies are wanted, but terminated

²⁸ Even if we *were* to aspire towards churning out stories through brute force, the time and effort spent doing so would surely then restrict our life-event experiences and even sensory experiences so severely that the concessions I gave to these two types of limitations would no longer reasonably hold.

²⁹ This is from a thought experiment by Tooley (1972).

because of extreme and tragic circumstances. Many late-term abortions, for instances, are conducted because of late diagnoses of severe fetal abnormalities which would cause the fetus to be incapable of surviving after birth, or be in severe pain and suffering. Yet, simplistic moral analyses of the issue (often by men) demonstrate ignorance of this fact (e.g. see Brooks 2018) that is well-known to those working in the field. Imagining this morally relevant possibility is theoretically possible without experience, of course. But in reality it is easy to miss this morally relevant possibility and imagine that women tend to undergo late-term abortions 'just because they can', which is a much less morally complex case to analyse. The point here is that such analyses are constructed on the basis of morally relevant features that one already has in mind, and what one is able to conceive of as morally relevant is limited by one's concrete experiences. This limitation greatly diminishes the value of the fact that imagination gives us access to contingent possibilities that concrete experiences cannot provide. While this is indeed an epistemic edge, being limited in experience means we will sometimes simply not know *what* to imagine, or fail to imagine cases that account for morally relevant features. So engaging in thought experiments without sufficient concrete experiences is useless.

One can rightfully point out that philosophers *have* indeed generated powerful thought experiments. This might indicate that philosophers have levels of concrete experience that are enough to gain moral knowledge, which goes against the idea that moral philosophers are not well-positioned to acquire expertise in

moral judgment. Or, this might indicate that philosophers have succeeded in generating powerful thought experiments *without* much concrete experience, which goes against the claim that concrete experience imposes important limitations on our ability to construct good thought experiments. But first, we should not assume that some of the most powerful thought experiments were not aided by moral expertise gained from the real world; Thomson, for instance, might be an example of this. Second, we should not assume that the vast amount of philosophical literature, and the vast number of what seem to be good thought experiments, have covered nearly enough of the moral truths we want to know, and that we demand of moral experts.

This point—about the wide scope of the moral truths we want to know—brings me to the last and most crucial way in which concrete experience limits one’s ability to construct thought experiments, and hence one’s ability to gain high levels of moral expertise. Recall that in order for an agent to count as a moral expert, she must be good at making *contingent* moral judgments in morally complicated situations. Lord and Milona acknowledge, as do I, that thought experiments allow one to acquire *some* amount of contingent moral knowledge. Yet the amount and sophistication of the contingent moral knowledge one can get via thought experiments is still limited by one’s concrete experiences. Without concrete experience, it is impossible to know how the actual world tends to exist and the combinations of morally relevant features that tend to crop up in specific moral sub-domains. Again, as an example, without experience it is easy to miss the contingent

fact that many late-term abortions are conducted because of severe fetal abnormalities that are hard to detect in earlier stages, and assume that women tend to get late-term abortions on a whim. Without being constantly exposed to the patterns of competing values that tend to contingently present themselves in the real world, it will be very difficult to cultivate the high level of sensitivity we demand from moral experts.

5.3 The challenge of thought experiments: round 2

There are other reasons why one might think that thought experiments alone cannot replace concrete experience.³⁰ I lack the space to consider them all. Instead, I shall end my discussion of thought experiments by delivering what was promised, and evaluating thought experiments as an option for the vast majority of the population who are not professional ethicists. Not everyone is as lucky as moral philosophers to be paid to do so. Given how much time and effort the average person can reasonably spend engaging in thought experiments, can the moral epistemic value of doing so replace the moral epistemic value of the time one spends effortlessly experiencing

³⁰ One tempting route to take is to claim that the affective reactions prompted by the imagination are somehow 'inferior' or less 'vivid' compared to the affective reactions prompted by concrete experience, in a way that makes them morally inferior. To my knowledge, no sentimentalists perceptualists have addressed this issue. For instance, neither Milona nor Lord seriously consider the possibility that imagination might produce less morally effective affective responses than concrete experiences. Lack of empirical evidence, however, makes it difficult to investigate this issue without tenuous appeals to personal phenomenology. There is one psychological study which suggests that emotional reactions to fiction are less intense than emotional responses to nonfiction (Cova et al., in prep). Without more studies, however, and without more studies specifically on imagination and *moral* emotion, it is still difficult to make this point.

the real world?

The answer is clearly no. Singer takes the average person's lack of time to engage in philosophy as a disadvantage: 'the moral philosopher can, if he wants, think full-time about moral issues, while most other people have some occupation to pursue which interferes with such reflection' (Singer 1972a, 117). Singer takes this to demonstrate that moral philosophers are much better-positioned than others to be moral experts. Yet, as we have seen, moral philosophising alone cannot generate the high levels of moral expertise we are after. Moreover, in ignoring the role of experience in generating moral expertise. what Singer takes to be the disadvantage of 'laypeople' is actually an advantage. One's very 'occupation to pursue which interferes with reflection' can provide the chance for agents to develop moral expertise in sub-domains relevant to those occupations. Singer underestimates the power of moral conflict and adjustment that 'the ordinary man' undergoes in his daily life, borne by concrete experiences with the real world.

One might point out that daily life *combined* with thought experiments can be more powerful than experience alone. And we need not limit our thought-experiment-like imaginative experiences to the sort of thought experiments philosophers entertain. Goldie, for instance, gives an account of how one can develop virtue by engaging in counterfactual thinking about one's past and present real-life experiences, which produces morally important affect (Goldie 2007, 354-356). And we do indeed engage in simple counterfactual thinking about our past

or prospective actions all the time. This is a crucial part of how we navigate the moral world. In deciding what we morally ought to do—or whether some action might be, for instance, cruel—we imagine *what might happen* if we did this, and what might happen if we did that instead. And this is certainly crucial for developing moral expertise.

However, this is no objection to the claim that experience is superior to ‘everyday thought experiments’ in generating moral expertise. This is because the very possibility of such counterfactual thinking depends on what we are experiencing in the concrete world. I would not be imagining ‘what might happen if I did this’ or ‘what might happen if I did that instead’ should I not be having the experience of having to make that choice. And my claim is not that engaging in such imagination cannot accelerate the process of developing moral expertise, or that being a moral expert does not involve such counterfactual thinking, or even that moral expertise is possible without engaging in such thinking. As I’ve noted, experience in itself is insufficient for generating moral expertise. The capacity to engage in such counterfactual thinking may well be necessary for this process. Yet engaging in this type of morally important imagination is only possible in virtue of one’s concrete experiences with the world.

5.4 The challenge of fictional narratives: round 1

Moving on from the challenge of thought experiments, I will move on to considering

fictional narratives as a means of gaining moral expertise, as compared to concrete experience.

In some ways, this might seem like a harder challenge to meet. Two big weaknesses of thought experiments are that on their own, they don't easily equip agents to handle morally complicated situations, and that they're bad at allowing agents to develop moral expertise in making valuable contingent moral judgments about the real world. But fictional narratives appear to fare much better in this respect. The moral power of fiction derives from its ability to paint complicated, compelling, and evocative pictures. Good fictional stories which are situated not too far from the real world put readers in a good position to gain knowledge of fairly sophisticated contingent moral truths about the actual world. Fictional stories that are more fantastical, like science fiction, can still be as morally powerful if their characters and plots bear relevant similarities to our experience of the real world. Moreover, the most morally powerful works of fiction allow us to step into the shoes of others, and simulate the experiences of people who are very different from us. Fictional works that allow us to successfully 'try on' the identities of others have a moral epistemic advantage over concrete experience, since we are able to simulate detailed experiences, in compelling ways, that we would otherwise not have had. Last, those in the 'narrative fiction' camp argue that fiction is a good way of cultivating one's very *moral sensitivity*—in other words, one's moral expertise—as opposed to merely being a good means to moral knowledge.

I am happy to grant that all this is true. In fact, I take all these points to count in my favour. As long as the moral power of fiction derives from allowing us to simulate experiences in detailed, realistic, and complicated ways, this means that its moral power is derivative upon the value of concrete experience in cultivating our moral sensitivities. The moral power of 'trying on' the identities of others through fiction, for instance, derives from the moral power of *being* those others. So, I argue, engaging with good fictional narratives will typically be a weaker way to develop high levels of moral expertise in specific moral sub-domains, as compared to having concrete experiences with the real world.

As with thought experiments, there are several ways to evaluate the moral power of fictional narratives. The first is to evaluate the power of reading good fiction under the constraints of ordinary life. In this case, it is quite clear that doing so is inferior to concrete experience. Temporarily stepping into some alternative world, for the period of time it takes to watch a movie or read a novel, is morally powerful. But surely this cannot replace the value of spending a *lifetime* actually *being* in those scenarios depicted, and actually *being* those characters described. If the characters in those movies were brought to life, whatever moral expertise we would have gained reading about or watching them would almost certainly pale in comparison to the moral expertise *they* would have. Watching all the episodes of *Star Trek: The Next Generation*, for instance, cannot possibly put me in the same position as Captain Jean-Luc Picard to become a moral expert at deciding how to interact with

alien species and cultures who have moral values very different from our own. Stepping into someone else's shoes for a short period of time cannot replace actually being that someone else.

5.5 The challenge of imagination: round 2

We can also evaluate the value of deliberately spending lots of time engaging with morally powerful fiction. At one extreme is someone who spends her entire life watching a movie—perhaps shot from a first-person perspective—depicting the entirety of someone else's life.³¹ While this is interesting, it is too unrealistic for us to usefully consider.

But there is at least one real life case study of extreme engagement with fiction that is interesting for us to investigate. This is the real life (and very unfortunate) case of the 'Angulo brothers'.³² For more than a decade, six brothers and one sister were locked inside a Manhattan apartment by their father, having no contact with the outside world. The vast majority of their time was spent watching movies nonstop, and it is estimated that they watched anywhere from two thousand to five

³¹ I leave it as an open question of fact as to whether this could replace the moral epistemic value of actually *being* that someone else. But note that this case actually blurs the line between 'concrete experience' and 'imagination'. Should the viewer spend her entire life literally doing nothing other than watching this movie, she has in some sense inhabited at least the visual and auditory faculties of the character in the movie. In some sense she has directly and concretely experienced, in a first-personal way, what it is like (visually and auditorily) to be that character.

³² Information about the Angulo brothers is from the American documentary film 'The Wolfpack' (Moselle, 2015), and the New York Times article "'The Wolfpack' Tells of One New York Apartment With Seven Children Locked Inside' (Barnes 2015).

thousand films. As Barnes reports, 'Quentin Tarantino, Christopher Nolan, David Lynch and Martin Scorsese gave them a window into the world ... and injected badly needed doses of creativity into their lonely, claustrophobic lives' (Barnes 2015). Much of their time was spent memorising and transcribing movie lines, creating costumes and props, and performing their favourite movie scenes for each other. Eventually, the six brothers escaped, and met Crystal Moselle, who eventually directed *The Wolfpack*, a 2015 documentary about their lives, starring the brothers themselves as the actors. Interestingly, the 'moral success' of the Angulo brothers' obsession with movies seems to be mixed. On one hand, the brothers have described facing difficulties connecting with people; one stated in the documentary that the hardest adjustment was 'having personal connections with other people', because '[their] own reality and our reality was so far removed'. Still, they appear to have turned out much better than one might expect after being locked up for more than a decade. For instance, they have described as being 'sensitive' and 'extremely likable' (Barnes 2015).

What we can conclude from the case of the Angulo brothers, of course, is limited. And it is still not descriptive of the majority of cases of those who deliberately spend lots of time engaging in good fiction. What can we say about them? I am willing to concede that it's much more likely that such agents, as compared to those who engage in extensive thought experiments, might be in a better position to acquire decent levels of moral expertise in specific sub-domains.

But this is provided that the moral scope of their fiction is sufficiently narrow and deep (i.e. they do not hop back and forth between narratives that represent experiences one might have in many different moral sub-domains), the fiction they engage with is morally powerful enough, and the way they affectively respond to fiction comes close to the way they would affectively respond to the real world. It's also important to note that fiction, depending on how it's written, can skew one's moral sensitivities in bad ways by not accurately representing certain descriptive patterns in the real world.³³ And these requirements so extremely stringent that it is quite unrealistic to think that the vast majority of individuals and the fictional works they choose could meet them. Even for the most avid and morally responsive reader, who chooses the 'morally best' narrative works that exist in the world, it is seriously unclear that this could ever match up to the value of concrete experience. And ultimately, most moral epistemic agents do not engage, and are not able to engage, with fiction this way. So, fictional narratives typically cannot compare to concrete experience when it comes to developing high levels of moral expertise.

³³ Interestingly, Crystal Moselle notes that 'the downside to all the movies [the Angulo brothers have watched] is that there are certain formulas to them. Real life is different. In real life, the girl doesn't always break your heart. The boys are still struggling to understand that.' (Barnes 2015).

6 Conclusion

To sum things up, I've argued for the main claim that experience with the concrete world puts one in an especially good position to gain expertise in moral judgment. But really, the scope of this paper has been quite large. In arguing for my main claim I've attempted to give answers to many other questions, and argued against various positions, in ways that will be quite contentious to many.

I've made a number of claims that will rub many moral philosophers up the wrong way. First, contrary to what they might wish to be true, I've argued that they aren't especially well-positioned to be experts in moral judgment after all. But moral philosophers need not be disheartened. The blow is softened by the fact that they count as good candidates for being experts in moral analysis and reasoning, which is another important form of moral expertise. Another argument I've made is that we should require moral experts to be good at gaining contingent moral knowledge about the actual world, and that the role of necessary moral truths ought to be de-emphasised on our picture of expertise in moral judgment. Yet moral philosophers must acknowledge that their own theoretical priorities should not shape what we demand of some everyday concept that the 'ordinary person' has a stake in. And nothing I've said means that there's no value in moral philosophers questing for necessary moral truths—this is obviously still a valuable activity in its own right. Another argument I've made is that philosophical thought experiments alone are

insufficient to generate the high levels of moral expertise in sophisticated contingent moral judgments that we demand from moral experts. However, this does not mean that thought experiments have no moral epistemic value—they do. Moreover, I've pointed out that thought experiments and counterfactual thinking, in combination with concrete experience, can be especially powerful. So nothing that I've said means that thought experiments are useless, or even not very useful.

This chapter also raises many avenues for further philosophical exploration. My proposal that we can undergo moral perceptual learning is a tentative one. I think it is an attractive proposal for moral perceptualists in general, but as it stands, it needs to be further fleshed out. Last, so far I have only been talking about the role of experience and moral sub-domains in quite general terms. But there is the possibility that, although experience is important for moral expertise in general, it will be much more important in certain moral sub-domains than others. Certain types of experiences with the concrete world might be much harder to replace with any kind of imagination than others. And this is something I will explore in Chapter 2: in the moral sub-domain of oppression, I will argue, it is so especially difficult to acquire moral expertise on oppression without having experienced being oppressed. What this means is that any social progress on the matter must often involve those in positions of power listening to the moral testimony of the powerless.

Chapter 2

Moral Expertise on Oppression

1 Introduction

A central tenet of standpoint epistemology is that members of oppressed groups have a richer perspective on the world than members of privileged groups. In other words, one's social privilege negatively correlates with one's epistemic privilege (Grasswick 2016). Although standpoint theorists make several more radical claims, one need not commit to all of them to see how being a member of an oppressed group can give one superior epistemic access to certain aspects of the world.

Consider the fact that women experience routine sexual harassment from men. This is something most women know in virtue of their own experiences. Yet, since sexual harassment tends to occur when women are alone, the phenomenon's existence and prevalence is impossible for most non-harassing men to know without relying on women's testimony. Moreover, even when men are provided with a full description of sexual harassment, the phenomenon is dissimilar enough from what

most men have experienced such that many of them find it difficult to grasp, for instance, the psychological toll it takes on its victims, the feeling of degradation experienced by women, and the relationship between instances of harassment taken in isolation and the routine, long-term sexism which women face in general. Consequently, many men incorrectly brush off phenomena like catcalling as being 'no big deal', and falsely believe that women complain about it for 'no good reason'. Many other oppressive phenomena are hidden from socially privileged knowers in this way, leading to false or absent beliefs about these phenomena.

Much has been said by standpoint epistemologists about how being a member of an oppressed group gives one access to more knowledge in general. In this paper, I aim to focus on the distinctively *moral* epistemic advantages that come with oppressed group membership. In particular, I'll argue that oppressed group members are often in a better position to become a *moral expert* with respect to the type of oppression experienced by that group, where moral expertise is defined as one's competence at forming moral knowledge in response to morally relevant features. For instance, women are in a better position to become moral experts with respect to the moral sub-domain of sexism, as opposed to men. They are in a better position to become competent at forming moral knowledge (on their own) about whether catcalling is morally wrong, the degree to which it is morally bad, whether a woman is blameworthy for not confronting her harassers, whether maternal leave policies are unfair, and so on. This stands in opposition to views held by

philosophers like Joseph Heath (2015), who hold that no such special epistemic advantages come with oppressed group membership, since ‘we all live in the same world’.

Just as this paper aims to add something new to discussions about the epistemic advantages of oppression, it also aims to add something new to the separate debate on moral expertise. Existing philosophers in this debate acknowledge the role of experience in producing moral expertise, and note that the diversity of human experience can, and often does, produce uneven levels of it throughout the human population.³⁴ So far, the discussion has been conducted on a general level. Philosophers have employed useful examples to illustrate these claims, and some of these have been about various forms of oppression.³⁵ But one can gain moral expertise with respect to many different issues, and the particular factors which tend to influence this process vary greatly depending on the moral sub-domain in question. There is a great deal to gain from honing in on particular moral sub-domains. So, these passing examples set the ground for an overdue, in-depth analysis of moral expertise in the specific moral sub-domain of oppression, which this paper aims to do.

Here's how I'll proceed. In Section 2, I'll outline the general account of moral

³⁴ See e.g. Driver 2006 & 2013, Jones 1999, Jones & Schroeter 2012 & 2018, McGrath 2011.

³⁵ Jones (1999), for example, notes that women can ‘come to have richer experiences with certain types of moral problems’, which can allow them to acquire ‘expertise about the kind of disvalue sexism is, and the often subtle forms it can take’ 65-66). In the nearby debate on moral understanding, Sliwa (2017) has noted that first-personal experiences of being a survivor of sexual assault, a victim of domestic violence, and a disabled person can greatly enable one's capacities for forming moral knowledge with respect to these issues (549-545).

expertise I'll be working with, discuss how experience works as a means towards becoming a moral expert, and discuss general factors which can help or harm this process. In Section 3, I'll apply this picture to the moral sub-domain of oppression, and argue for the main claim that a member of an oppressed group is often in a better position to become a moral expert with respect to the type of oppression experienced by that group. In my conclusion, I'll briefly discuss some of the specific factors which can prevent oppressed group members from becoming moral experts on their own oppression.

2 Moral expertise

2.1 What is moral expertise?

In Chapter 1, I've focused on the following definitions of moral expertise and experts:

Moral expertise: Moral expertise within moral sub-domain D is an agent's competence at forming moral knowledge in response to morally relevant features, within D .

Moral expert: A moral expert in some moral sub-domain D is an agent whose moral expertise allows her to reliably form moral knowledge within D .

For instance, one can be a moral expert in the moral sub-domain of friendship, by

being sufficiently competent at forming moral knowledge with respect to a broad range of questions about friendship, and by doing so in response to the morally relevant features which ground those moral truths. This involves knowing whether a certain act counts as a betrayal, whether one is blameworthy for cutting off a difficult friend, and so on.

Here are some important things to note. First, the type of moral expert I'm interested in is the moral *epistemic* one, even though there are obviously many other senses in which one can be a 'moral expert'.³⁶ Second, being competent at forming moral knowledge isn't sufficient for one to count as a moral expert; one must be competent at doing so in response to the *morally relevant features* of the situation. This eliminates the possibility that one can be a moral expert solely by being competent at gaining moral knowledge via testimony.³⁷ Third, a moral expert needn't be capable of communicating, or even articulating to herself, the moral reasons which ground her knowledge. It's possible for one to reliably form moral knowledge in response to morally relevant features without having these capabilities.

Fourth, when it comes to moral sub-domains which are individuated by real-

³⁶ In addition to being competent at forming moral knowledge, one can also be competent at performing right actions, or engaging in moral reasons. These strands of moral expertise frequently come apart, however, and having one form doesn't guarantee having the others. For instance, excelling in moral reasoning, as many philosophers do, is no guarantee that one will excel at forming knowledge; the diversity of positions within normative ethics means that the vast majority of moral philosophers have false moral beliefs with respect to these questions.

³⁷ The morally relevant features which make it the case that it's impermissible for me to betray my friend, for instance, may include her potential hurt feelings, her loyalty to me in the past, and that she is my friend. I may be good at reliably picking out others who can provide me with moral knowledge about whether I'm permitted to betray her, without responding to any of the morally relevant features. In such a case, what I'm responding to instead are the epistemic features of the situation which weigh in favour of trusting someone else's testimony.

life phenomena, it isn't sufficient or even necessary for moral experts to be competent at forming moral knowledge about necessary moral truths and abstract moral principles. They must be competent at forming moral knowledge of highly contingent, everyday moral truths. It's one thing to know, as many of us do, that 'betraying a friend is morally bad', and that 'one ought to weigh costs and benefits when deciding whether to lie to a friend'. But it's quite another thing to know, for instance, whether to lie to a very depressed friend with extremely low self-esteem that her bad drawings, which she's spent days slaving over, are good enough to post online. Setting aside the interests of moral philosophers, the everyday person is typically fundamentally interested in moral knowledge of the latter type. And quite plausibly, one need not even have knowledge of the former, more abstract principles in order to know what to do in complex cases like these.

2.2 Context and combinations of morally relevant features

In previous work, I've also argued that extended experience with the real world is typically crucial, and irreplaceable, for gaining high levels of moral expertise in many moral sub-domains. But this process relies heavily on the *context* of an agent's experiences, and the *combinations* of morally relevant features she encounters. So I'll first explain these two things and what they have to do with moral expertise. The main point here is that moral expertise within some moral sub-domain typically involves being responsive to characteristic combinations of morally relevant features, when they are present in specific contexts which are distinguished by non-

moral features.

Let's begin with *context*. We often pick up on morally relevant features in some contexts but are insensitive to those same features in others, where contexts are distinguished from each other by morally irrelevant (or non-moral) features. For instance, almost everyone is highly sensitive to the morally relevant feature of physical pain when it's instantiated in human beings, but fewer people are sensitive to it when it's instantiated in non-human animals. Consequently, even when there are no morally relevant differences between cases of human and non-human animal pain, people often draw different conclusions about each type of case. While context sensitivity clearly produces pernicious biases, it's also a plausible byproduct of experience as a method of producing moral expertise. Since the type of experience that confers moral expertise tends to repeatedly expose us to characteristic combinations of moral and non-moral features, it's no surprise that moral experts will be sensitive to some morally relevant features when they're accompanied by certain non-moral features, and insensitive to those same morally relevant features when those non-moral features are lacking.³⁸

Now, here's the significance of being responsive to *combinations* of multiple morally relevant features. Most of us plausibly count as being moral experts in the

³⁸ For instance, repeatedly experiencing and witnessing racism against the Asian community can allow one (if all goes epistemically well) to become a moral expert in anti-Asian racism. It's possible, however, that the agent's superior sensitivity towards the morally relevant features of racism (e.g. degradation, mental suffering, objectification) remains dependent on the presence of non-moral features that have typically accompanied her experiences (e.g. the target being Asian, rather than Black), such that she remains incompetent at forming moral knowledge in the sub-domain of anti-Black racism.

moral sub-domain of ‘relatively simple questions about human physical pain’, in virtue of our moral sensitivity to physical pain when instantiated in humans. But moral expertise in more substantive moral sub-domains, such as animal ethics, requires competence at forming moral knowledge in response to *multiple* morally relevant features, which tend to come in certain characteristic combinations. For example, a moral expert in animal ethics must be morally sensitive towards physical pain in non-human animals, human interests like staying out of hunger, the value of human cultural food practices, and so on. And it isn’t enough for one to be sensitive to all of these in isolation; one must also be sensitive to the combination of all these features in the right way. Many of us have a decent amount of moral sensitivity to all these matters alone, but are not very good at weighing the interests of humans and non-human animals when they conflict.

2.3 Experience as a path towards moral expertise

Now, here’s a simplified, stepwise account of how an agent’s extended experiences with the world can help her become a moral expert in some moral sub-domain. Experience, as I’ve argued elsewhere, is a superior means towards becoming a moral expert which typically can’t be replaced by mere imagination. But importantly, I’m not arguing that experience is necessary in principle for becoming a moral expert; I’ll allow for the theoretical (but unlikely) possibility that someone can become a moral expert without it. Also, the claim isn’t that experience guarantees moral expertise; as I’ll discuss shortly, many obstacles can stand in the way. All that any prior step does

is to put one in a *position* to reach the next step.

Consider some moral sub-domain D , which is characterised by certain combinations of morally relevant features (m_1, m_2, m_3) , in context C involving certain combinations of non-moral features (n_1, n_2, n_3) :

- Step 1: An agent encounters (m_1, m_2, m_3) in the concrete world in context C , by experiencing it for herself, or witnessing it second-hand.
- Step 2: The agent has some recognition of (m_1, m_2, m_3) as morally significant to some degree.
- Step 3: Owing to this encounter, or over the course of subsequent encounters with (m_1, m_2, m_3) in C , the agent becomes sufficiently morally sensitised such that in future cases where (m_1, m_2, m_3) is present in C , she can reliably detect and recognise (m_1, m_2, m_3) as morally significant to some degree.
- Step 4: The agent becomes competent at using this moral sensitivity to reliably form moral knowledge in moral sub-domain D . In other words, she becomes a moral expert in D .

Here are some important things to note. First, this process isn't necessarily tied to any particular method of responding morally to the world. It's a pluralistic account of the faculties which may be involved in moral expertise. One's 'recognition' of the

moral significance of some morally relevant feature, and one's 'moral sensitivity' which may be used to form moral knowledge, may be realised by one or more different faculties including or cognitive, affective, and conative faculties. Consider the example of physical pain. We typically respond to this with negative emotions, and a desire for it to stop. For almost all human beings, experiencing physical pain first-hand and responding to it in these ways will be our first insight into its moral significance. But we often also respond to physical pain in more cognitive ways, by forming explicit beliefs like 'the physical pain of a drowning child ought to be stopped', and by using moral reasoning to form moral knowledge. This accommodates the fact that agents have different cognitive, affective, and conative faculties, and that there are multiple ways in which we respond morally to the world to form moral knowledge.³⁹

Second, Steps 2 and 3 of the process require agents to recognise the moral significance of morally relevant features to at least *some* degree. But this recognition need not be maximal, such that we e.g. make explicit moral judgments like 'The pain I'm experiencing here is morally bad', or feel intense moral anger at witnessing one's colleague being sexually harassed. An agent who merely suspects that something is

³⁹ In this respect, my account of moral expertise bears similarities to Sliwa's (2017) account of moral understanding, which she defines as 'the ability to acquire knowledge' (546). According to Sliwa, 'the capacity of moral understanding can be multiply realised. Plausibly there are many distinct faculties and cognitive mechanisms by which we acquire moral knowledge: perception, imagination, intuition, our affective responses, and moral reasoning can all be sources of moral knowledge. ... [M]oral understanding is realised by a set of different faculties and cognitive abilities. Agents can have these cognitive abilities to different degrees. And so, what exactly grounds the capacity of moral understanding may vary from agent to agent' (548). Sliwa, however, argues that one's ability to form moral knowledge purely via testimony can count as moral understanding; my account does not allow for testimonial competence to count as moral expertise.

morally fishy about the way her boss is treating her colleague, or feels a little upset that something seems 'off', clearly has at least some insight into the moral significance of the sexual harassment she witnesses.

2.4 Factors which can help and harm experience

Numerous factors can influence how likely an agent is to progress from one step to the next. Here are just some examples.

First, although moral sensitivity isn't necessarily tied to any particular faculty, it's likely that one's affective and conative faculties will be especially valuable in many cases. Consider the following example. The Notre Dame cathedral is currently engulfed in flames. Given all the information I have, I form the piece of evaluative knowledge that very many people are suffering a great loss, and that something very valuable is being damaged. Despite this, I myself have no affective or conative responses towards the event, since I feel no personal connection towards the cathedral, or for that matter, historical and cultural artefacts in general. When my friend says, 'A lot of people are concerned about Notre Dame burning down, but no one is talking about where Quasimodo is going to live now,' I burst out laughing at his joke, and proceed to post it on Facebook for all to see. My Parisian friend chides me for propagating the joke to those in mourning while the cathedral is still burning. But I fail to see why I should take my post down, and believe that she's overreacting. Here, my lack of affective and conative sensitivities towards Notre Dame and other cultural artefacts partly explains why I've found the joke so funny. It also explains

why I fail to be competent at knowing that repeating it to those in mourning as the cathedral burns is at least morally suberogatory or offensive, despite having a sufficient amount of information to reason my way to this conclusion. It will be difficult for me to become a moral expert in how to respond in the wake of tragedies, at least without any extra help. In general, lacking fitting affective and conative responses will often, though not always, be an obstacle in the path towards becoming a moral expert.

Second, one's background beliefs can influence one's likelihood of progressing through the steps, and this can itself be influenced by factors such as one's upbringing. Consider Emily, who reliably experiences negative moral emotions when she witnesses non-human animals suffer. After watching the movie *Oldboy*, she experiences disgust and horror when the protagonist eats a live, wriggling octopus. Emily knows that cephalopods are capable of experiencing large amounts of physical pain. Her upbringing, however, has conditioned her to believe that eating live seafood is morally permissible as an integral part of her culture's eating practices. So, she ultimately gives a moral pass to the protagonist (and the actor Choi Min-sik). Here, despite Emily having a large amount of affective insight into the moral significance of physical pain in non-human animals, her other beliefs stop her from endorsing her moral emotions, and prevent her from reliably forming moral knowledge with respect to certain types of moral questions. She fails to progress from Step 3 to 4, which prevents her from becoming a moral expert in the sub-

domain of animal ethics.

Third, in many cases, having first-personal experience of certain moral phenomena seems to boost the likelihood that one will recognise, and to a greater degree, the moral significance of the morally relevant features present. Consider the question of whether someone with an eating disorder is at fault for failing to ‘snap out’ of her condition. One morally relevant consideration here is that her emotions, desires, beliefs, and visual perception of her own body are subject to a foreign, manipulative influence. Someone who’s experienced her agency being undermined in this way is in a better position to recognise its moral significance, and is also in a better position to recognise this to a greater degree. On the flip side, the experience of having an eating disorder is sufficiently distinct and dissimilar to most other ‘normal’ human experiences, such that someone who hasn’t experienced this illness will have a moral epistemic handicap, and will be less competent at knowing on their own that victims of eating disorders are not to blame for their own condition.

Fourth, when one’s experience isn’t first-personal, witnessing some wrong being done to others we *care* about boosts the likelihood that we will recognise the moral significance of what’s transpired. Someone who sees her mother being mocked for her ‘funny Asian accent’, for instance, is much more likely to respond with a greater amount of indignation than another bystander who has no relationship to the victim of this racist encounter.

Fifth, having access to others with moral expertise can also boost the likelihood

of one becoming a moral expert. Consulting a therapist, for instance, often helps those with eating disorders to make sense of their experiences, recognise that they have a mental illness, and recognise that their self-blame is irrational. And this isn't merely a matter of the victim believing her therapist's testimony that she isn't to blame. Undergoing therapy can help her recognise *for herself* the morally relevant features of her condition, and put her in an excellent position to develop her own moral sensitivities.

Sixth, having access to the relevant communities is valuable in a similar way. It's common for victims of eating disorders and other mental illnesses to seek out online or offline communities of other victims. Exchanging experiences, engaging in collective reasoning, and listening to community members with high levels of moral expertise, can put an agent in an excellent position to develop her own moral sensitivities, and become a moral expert in the sub-domain of eating disorders or mental illness.

The factors which can influence one's likelihood of becoming a moral expert are numerous, and the list above isn't exhaustive by any means. And the particular factors which tend to be helpful or harmful vary from sub-domain to sub-domain. In the next section, I'll discuss the particular factors which tend to help or harm one's likelihood of becoming a moral expert in sub-domains of oppression.

3 Moral Expertise on Oppression

In this section, I'll argue for one main claim. The main claim is that a member of an oppressed group is often in a better position to become a moral expert with respect to the type of oppression experienced by that group. For instance, women are in a better position than men to become moral experts in the moral sub-domain of sexism, disabled people are in a better position than the able-bodied to become moral experts in ableism, fat people are in a better position to become moral experts in fat-shaming than those who've always been thin, and so on.

I'll begin by introducing the nearby concept of the 'standpoint', since my argument will frequently draw support from what standpoint epistemologists have to say (3.1). Next, I'll show how the experiences that come with being in an oppressed group can help one become a moral expert in the moral sub-domain of that type of oppression. This is because oppressed group members are more likely to be repeatedly exposed to oppressive phenomena (3.2), and also because they are more likely to recognise the morally relevant features involved as being morally significant, become sensitised to these features in the appropriate contexts, and become competent at using moral sensitivities to form moral knowledge (3.3). Last, I'll respond to the objection that just as the interests of the privileged are likely to lead them epistemically astray, the interests of the oppressed are also likely to do the same (3.4).

3.1 Standpoint epistemology and moral expertise

Nearby to the notion of becoming a moral expert on oppression is what standpoint epistemologists have called achieving the 'standpoint'. I'll introduce the concept here, and say what this has to do with moral expertise on oppression, since my subsequent discussion will draw support from some of the claims made by standpoint epistemologists.

Nancy Hartsock (1983) defines a standpoint as a 'stance' that reveals to socially underprivileged knowers the shared nature of their experiences of oppression, and the systematic structure of power relations. Standpoint theorists hold that one's social location—i.e. where one exists along axes of oppression— influences one's epistemic privileged, and that the oppressed are much better positioned to achieve a standpoint than dominant groups.

While there are important differences between my claims and the those that standpoint epistemologists make, I'll first focus on the similarities. First, both occupying a standpoint and being a moral expert on oppression involve having some type of richer epistemic perspective on the type of oppression faced by the marginalised group in question. Second, both standpoint theorists and I claim that being a member of an oppressed group puts on in a better position to have this richer epistemic perspective on oppression. Third, occupying a standpoint and being a moral expert on oppression are both products of an achievement, and aren't automatically guaranteed by being a member of an oppressed group. Fourth, and

relatedly, the likelihood of achieving a standpoint and becoming a moral expert on oppression are influenced by similar factors.

Given these similarities, the following discussion of moral expertise on oppression will frequently draw on support from standpoint epistemologists. But readers should not assume that standpoint epistemologists will agree with all of my claims. In fact, many will disagree with some of my key meta-epistemic commitments and philosophical methodology.⁴⁰ On the flip side, readers should also not assume that I agree with all that standpoint epistemologists have to say, beyond the similarities I've outlined above. Still, even where there is technical disagreement between standpoint theorists and I, my account has much to gain from taking many of their claims metaphorically.

3.2 Being exposed to oppression

In the previous section, I outlined how experience can be a powerful means of gaining moral expertise. Here, I'll discuss how the experiences that come with being in an oppressed group can help one become a moral expert in the moral sub-domain of the oppression experienced by that group.

I'll begin by addressing Step 1 and part of Step 3 of the experience-expertise process: where an agent encounters, and continues to encounter, combinations of morally relevant features in the world in the relevant contexts. Oppressed agents are

⁴⁰ For example, some will reject the notion that there is a class of knowledge which is distinctively 'moral'. Many of them are also likely to reject my account of moral expertise on the grounds that it is too essentially individualistic.

much more likely to encounter and be repeatedly exposed to the oppressive phenomena which are constituted by these features and contexts.

As I noted in the introduction, society is often structured in ways that many oppressive phenomena are literally obscured from the view of the non-oppressed. Jaggar notes that the class position of the privileged 'insulates them from the suffering of the oppressed', such that they 'encounter little in their daily lives that conflicts' with the belief that 'current organisation of society [is] basically satisfactory' (2004, 56-57). Consider again the case of catcalling; since most non-catcalling men are rarely privy to the phenomenon, they simply do not encounter the particular cases in which the morally relevant features related to catcalling are instantiated. Let's assume that some wrong-making features of cat-calling are the unwarranted shame it tends to cause women, the objectification that occurs, the fear of one's physical safety that it causes, and its contribution to the constant mental toll that sexism takes on women in general. Even though many men have experienced some of these morally relevant features (e.g. unwarranted shame, fear of safety, mental tolls of various kinds), they are less likely to have experienced some of them (e.g. objectification), and they are also much less likely to have experienced the particular combination of these morally relevant features.

An important factor which can influence one's likelihood of encountering oppression is where one lives. In many societies, this often correlates with one's social location, defined in terms of class and race. Those who live in higher-income

neighbourhoods are much less likely to witness the difficulties faced by those trapped in poverty cycles. And those who live in predominantly white neighbourhoods (which are also often higher-income neighbourhoods) are much less likely to witness Black people being treated differently by the police than they are. One's social circle and community is also an important influence. Once again, this is often determined by factors like income, race, sexual orientation, and so on. Someone with no transgender friends, for example, is less likely to witness instances of discrimination faced by the transgender community. They are less likely to witness someone struggling to decide which bathroom to use in a nightclub, out of fear of being kicked out for using the 'wrong one'.

Consider also the fact that inequality is one morally relevant feature of oppression. One can only fully encounter inequality by bearing witness to both oppression and privilege, and for many types of oppression, marginalised agents will be much more likely to enjoy this dual perspective. Patricia Hill Collins discusses the example of how many Black women caretakers have special insight of this kind:

'Afro-American women have long been privy to some of the most intimate secrets of white society. Countless numbers of Black women have ridden buses to their white "families", where they not only cooked, cleaned, and executed other domestic duties, but where they also nurtured their "other" children. ... These women have seen white elites, both actual and aspiring,

from perspectives largely obscured ... from these groups themselves. (Collins 2004, 203).

One thing we can take away from this quote is that members of oppressed groups, such as Black women, often have jobs that allow them to see how privileged individuals live in comparison to themselves. This is often true of women, ethnic minorities, and lower-income individuals, who are disproportionately tasked with jobs which require them to service the elite. But the privileged typically have no reason to take a converse peek into the lives of the oppressed. In this sense, the comparatively privileged nature of their own circumstances is 'obscured from themselves'.

This asymmetrical exposure to inequality is not limited to cases of servicing the wealthy. Disabled people, for instance, bear witness to the relative privilege of able-bodied people every time they leave their homes to interact with the outside world. They witness the relative ease with which able-bodied people navigate their physical environments, and observe able-bodied people being able to enter venues and use facilities which they themselves cannot. Able-bodied people, in contrast, are much less likely to bear witness to the struggles of disabled people, since the latter form a minority of the population.

3.3 Recognising moral significance, being morally sensitised, and becoming a moral expert

Steps 2, 3, and 4 of the experience-expertise process involve the agent recognising morally relevant features as being morally significant to some degree, becoming morally sensitised to those features in the appropriate contexts, and becoming competent at using her moral sensitivities to form moral knowledge in the relevant moral sub-domain. Once again, oppressed agents are in a much better position to make this progress, with respect to the sub-domains of their own oppression.

As I've noted in the previous section, first-personal experience of being wronged and harmed often puts us in a better position to recognise the moral significance of what's happened. This is true of the wrongs and harms which contribute to oppression. Someone who's on the receiving end of a fat-shaming encounter, for instance, is typically more likely to have a negative moral reaction towards the event than a thin-bodied witness to this event. Moreover, as I've also previously noted, witnessing wrongs and harms being done to those we care about also boosts the likelihood that we'll recognise the moral significance of what's being done. In moral sub-domains such as racism, marginalised agents will often witness these oppressive wrongs and harms being done to family members they care about. Recall the example of someone witnessing his mother being mocked for her Asian accent; he's more likely than a white stranger to feel indignant at what's happened. This point extends from family to friends as well, at least when social groups cluster

along lines of oppression and privilege. The friend who witnesses her transgendered friend struggling to decide which bathroom to use, for instance, is much more likely to feel sadness and anger at her friend's predicament than a stranger listening to their conversation. And as I've noted in the previous section, affect is often an especially powerful tool when it comes to moral recognition, and developing moral sensitivity. In many cases of oppression, an oppressed person being wronged or harmed, or witnessing someone they care about being wronged or harmed, will respond with greater amounts of affect than someone who is not a member of that oppressed group.

Experience which is likely to trigger moral recognition need not always be of wrongs and harms in order for one to be in a better position to make certain moral judgments about oppression. For example, Jaggar (2004) notes that 'it is only from the standpoint of women that household labour becomes visible as work rather than a labour of love'. Having to do household labour is not a wrong or harm in itself. It does, however, put women who are relegated this role in a better position to recognise that their love for their family members does not generate a moral duty for them to take on household responsibilities for free. It also puts them in a better position to recognise that women are being unfairly treated if their household labour isn't treated as work, when their husbands' non-household jobs are.

The interests of oppressed agents can also push them in the direction of recognising the moral significance of oppressive phenomena, becoming morally

sensitised to it, and eventually becoming competent at forming moral knowledge about it. On the other hand, non-marginalised groups often have interests in maintaining is not as prevalent or bad as it really is, or that what they witness is not really oppressive. This point about interests is a popular sentiment amongst standpoint epistemologists. Jaggar, for example, says the following about the purported 'knowledge' perpetuated by the dominant class:

'In a society where the production of [purported] knowledge is controlled by a certain class, the [purported] knowledge produced will reflect the interests and values of that class. ... Because the ruling class has an interest in concealing the ways in which it dominates and exploits the rest of the population, the interpretation of reality that it presents will be distorted in characteristic ways. In particular, the suffering of the subordinate classes will be ignored, redescribed as enjoyment or justified as freely chosen, deserved, or inevitable.

'Because their class position insulates them from the suffering of the oppressed, many members of the ruling class are likely to be convinced by their own ideology; either they fail to perceive the suffering of the oppressed or they believe that it is freely chosen, deserved, or inevitable. They experience the current organisation of society as basically satisfactory and so they accept the interpretation of reality that justifies that system of organisation. They encounter little in their daily lives that conflicts with that

interpretation.

‘Oppressed groups, by contrast, suffer directly from the system that oppresses them. Sometimes the ruling ideology succeeds in duping them into partial denial of their pain or into accepting it temporarily, but the pervasiveness, intensity, and relentlessness of their suffering constantly push oppressed groups toward a realisation that something is wrong with the prevailing social order.

‘The standpoint of the oppressed is not just different from that of the ruling class; it is also epistemologically advantageous.’ (Jaggar 2004, 56-57)

Similar sentiments also exist in the Marxist literature. Grasswick describes the Marxist materialism developed by Georg Lukács, who describes the epistemic asymmetry as one of 'understanding':

‘One’s social position with respect to material labour is inversely related to one’s epistemic position. ... As the privileged class, the capitalists have a motivation to maintain the status quo, and this interest interferes in their ability to understand the exploitation of the working class upon which their capitalist privilege depends. The working class, however, as the socially underprivileged, can achieve a richer understanding of social relations; they ... have a motivation to understand the true nature of the exploitation to which they are subject (in order to be able to end the exploitation). ... Thus, their position as socially underprivileged affords them the possibility of an

epistemic privilege'. (Grasswick 2016, 14-15).

Whether they recognise it or not, members of privileged groups often have interests in maintaining false beliefs about oppression, since maintaining the status quo benefits them; the converse is true about members of oppressed groups. Not only can this affect the production of knowledge and moral expertise, it can plausibly also affect the likelihood of one being affectively, conatively, or cognitively responsive to oppression in the right ways.

3.4 The objectivity objection

At this point, a natural worry arises about the influence of interests on how likely an agent is to become a moral expert on oppression. What I've just claimed is that the interests of the marginalised make them more likely to become moral experts on their own oppression, whereas the interests of the privileged make them less likely to do the same. To put things another way, I've claimed that the moral perspective of the marginalised is more *objective* than that of the privileged, when it comes to matters of oppression. To some, this might seem too convenient. In general, when we think about two agents with contrasting interests, we tend to think that each of their perspectives are likely to be distorted in opposite directions. Neither party is in a good position to be objective, and both are likely to make just as many false judgments about the matter. Call this the 'objectivity objection'.

With respect to moral expertise on oppression, this objection is largely

unfounded. First, note that there are obvious historical counter-examples to this claim. One such case is that of Afro-American slavery. White slaveholders had interests in maintaining that slavery was morally permissible, while Black slaves had interests in maintaining that it was not. Yet, it's clearly false that each set of interests was equally epistemically distorting. Moreover, the interests of Black slaves were not just non-distorting, or less distorting—they were positively truth-tracking. This leads to a broader point, which is that the objectivity objection rests on the false assumption that interests must always be distorting, and can never be truth-tracking. Louise Antony (1993) acknowledges that although 'biases' are an inevitable component of knowing, it's ultimately an empirical question as to which biases are epistemically good or bad.

So, when is there likely to be an epistemic asymmetry in the value of contrasting biases? One simple answer is that this tends to happen when one party has interests in maintaining some deception or falsehood, while the other has interests in uncovering the deception. Both Jaggar and Lukács appeal to this idea. But this may be unsatisfying to some since often, the very question at stake is whether or not some deception is at play. Although this is an epistemic worry, since in these cases there is a fact of the matter, it's still worth thinking about how we can identify which party is which. One question we can ask is whether there seems to be an existing external force which has a clear interest in both parties believing something irrespective of its truth, and which also has the power to influence one or

both parties into doing so. There are clear non-moral examples of this. Consider the powerful influence of the American pharmaceutical industry, which gives us reason to believe that the financial interests of doctors with respect to prescribing certain medications are epistemically distorting, whereas the interests of patients in maintaining their health are truth-tracking.⁴¹

Many cases of oppression will be like this. The powerful, external distorting force at play is what philosophers sometimes call the dominant 'ideology'. Ideology, roughly, is the set of ideas, ways of thinking, and beliefs predominant in society. Ideology which is oppressive is influenced by at least two things. The first is the group of privileged individuals who are aware of the truths about oppression, but perpetrate falsehoods about it anyway, such as owners of sweatshops who know their workers are being perniciously exploited. The second influence on ideology is the set of structural forces either produced by these agents in the know, or produced non-agentially by causes such as the natural tendency for society to be structured in certain ways without intervention. One example of the latter is the tendency for physical environments to be built for the convenience of the able-bodied, since they

⁴¹ The privatised pharmaceutical industry in America has financial interests in getting both doctors and patients to believe that their medication is the safest and most effective solution to most health problems, irrespective of whether this is the case. The result of this is that doctors have interests in maintaining that this is true, since they often receive financial incentives from pharmaceutical companies to prescribe those medications. Moreover, evidence suggests that these incentives are often powerful enough to outweigh the goodwill or medical expertise of doctors; research shows, for instance, that doctors who receive more money from opioid manufacturers are significantly more likely to prescribe opioids to their patients. Patients, on the other hand, should be motivated to find out what is in fact the truth, since it's in their best interests to follow the course of action which would best improve their health. The result of this is that we have reason to think that the interests of doctors with respect to the matter are more epistemically distorting than the interests of patients, which are much more truth-tracking.

form the majority of the population. Another example is the tendency for women's careers rather than men's to take a hit after childbirth, since without intervention it often makes more financial sense for women rather than men to put their careers on hold after childbirth.

The 'interest' that oppressive ideology has in perpetrating itself comes both from the interests of malicious, privileged individuals, as well as the inertia society must overcome in order to structure itself in opposition to the contents of that ideology. Once we examine the contents of the dominant ideology, it is easy to identify some cases in which privileged individuals have interests in maintaining the status quo, and cases in which oppressed individuals have interests in uncovering the truth. Given the natural tendency for physical environments to be built for the able-bodied, and the extra resources it would take to accommodate disabled people, we have good reason to think that the moral sensitivities of the able-bodied are likely to be distorted when it comes to moral judgments about ableism, and that those of disabled people are likely to be truth-tracking. Given that sweatshop owners have interests in maintaining that their employees' working conditions are morally permissible, irrespective of the truth, we have good reason to think that the moral sensitivities of many consumers are likely to be distorted due to their interests in keeping their cost of living low, and that the sensitivities of sweatshop workers are likely to be truth-tracking.

This type of reasoning can serve as a good epistemic guide to determine the

epistemic value of biases in *some* cases of oppression. But we must be careful not to think that it will help us uncover *all* of them. It's unclear, for instance, as to how forms of oppression such as fat-shaming fit into this picture (especially since the powerful fast food industry has interests in maintaining that there are no moral problems with obesity).

Ultimately, though, this epistemic worry shouldn't keep us from thinking that there are many more truth-tracking biases than we can readily identify. Standpoint epistemologists have had much more to say about the false assumption that epistemic objectivity necessarily involves 'neutrality' or a lack of interests. Sandra Harding notes that examining cases of oppression gives rise to a 'critique of conventional epistemic standards', and reason to think that 'what Donna Haraway dubbed 'the God Trick'—the traditional epistemic view that knowledge is only achieved by adopting a disinterested, impartial view from nowhere—is unachievable, for knowledge is always from somewhere' (Harding 2004, 93). Alison Wylie also argues that such cases 'explicitly undermine the assumption that objective epistemic agents are non-specifically located, and that they are neutral and disinterested with respect to the subject of their inquiry' (Wylie 2004, 345-356).

4 Conclusion

My main claim has been that a member of an oppressed group is often in a better position to become a moral expert with respect to the type of oppression experienced by that group. So far, I've only discussed the epistemic advantages that come with being an oppressed group member. But being oppressed is no guarantee that one will become a moral expert on the type of oppression one experiences, and I haven't discussed the specific factors which can prevent oppressed agents from becoming moral experts on their own oppression. Although I lack the space for a full discussion of the matter, here is a brief one.

The biggest obstacle, with respect to many types of oppression, will be oppression that marginalised agents have internalised. Here is Jaggar:

'The daily experience of oppressed groups provides them with an immediate awareness of their own suffering but they do not perceive immediately the underlying causes of their own suffering nor even necessarily perceive it as oppression. Their understanding is obscured both by the prevailing ideology and by the very structure of their lives.

'In addition to the mystifications created by the dominant ideology and by the structure of our lives, ... women face another obstacle ... The obstacle is the typically feminine set of attitudes and modes of perception that have been

imposed on women in a male-dominated society. ... While women's experience of subordination puts them in a uniquely advantageous position for reinterpreting reality, it also imposes on them certain psychological difficulties which must themselves be the focus of a self-conscious struggle.

'Simply to be a woman, then, is not sufficient to guarantee a clear understanding of the world.' (Jaggar 2004, 60-61)

Many survivors of sexual assault, for example, have been conditioned by prevailing societal attitudes to feel shame and self-blame. Many women incorrectly endorse these emotions, and believe that they are in fact to blame for their own sexual assault. As Jaggar notes, a 'self-conscious struggle' against such emotions is necessary for women to recognise that these emotions are a misleading product of their upbringing and environment.

Another obstacle which often hinders marginalised agents is the phenomenon of hermeneutical injustice. Hermeneutical injustice, as defined by Fricker (2007), occurs when there is a deficit in our shared tools of social interpretation (i.e. the collective hermeneutical resource), such that marginalised social groups are at a disadvantage in making sense of their distinctive and important experiences. Before the term 'sexual harassment' was coined, for instance, many women had an impoverished, scattered understanding of being sexually harassed, and were hindered from forming moral knowledge about the wrong being done to them. Having thick ethical concepts like 'sexual harassment' greatly increases one's ability

to form certain types of moral judgments. More recent sexist moral phenomena that have been collectively identified include 'mansplaining', 'he-peating', and 'correctile dysfunction'; having names for these phenomena have allowed many women to make sense of their experiences, identify them more confidently in the future, confidently and justifiably recognise that they are being wronged when they experience these phenomena, and form moral knowledge. But there are undoubtedly a great number of oppressive phenomena that have yet to be given a name, and this provides an obstacle towards moral expertise with respect to various types of oppression.

Related to these two points is the importance of community in influencing how likely one is to become a moral expert. In Section 1, I noted that having access to communities comprised of others with similar experiences can boost the likelihood of one becoming a moral expert. Whether or not such communities exist for oppressed agents can greatly harm or hinder their chances. Exchanging experiencing and engaging in collective reasoning can boost the likelihood that oppressed agents will engage in Jaggar's 'self-conscious struggle' against harmful psychology, and it can also boost the likelihood that oppressive phenomena will be given names, and that hermeneutical justice will prevail.

To wrap things up, I've applied an experience-based picture of moral expertise to the moral phenomena of oppression. I've discussed how being a member of a marginalised group puts one in a better position to become a moral

expert with respect to the moral sub-domain of the oppression experienced by that group, as opposed to non-members of that group. But I've also noted how being oppressed doesn't guarantee that one will become such a moral expert, and briefly listed just some obstacles that can stand in the way.

Importantly, what's been said in this paper is by no means intended to be exhaustive, or universal to all types of oppression. There are many forms of oppression, and each of them comes with unique features and complications. This cursory analysis of oppressions in general has highlighted the richness and complexity of how moral expertise operates in the real world, and zooming into specific forms of oppression even further will reveal even more.

Last, this discussion has important epistemic implications for whether, and under what circumstances, we should defer to members of marginalised groups when they give moral testimony about their own oppression. In certain circumstances, privileged individuals will have strong reason to do this, but as the discussion of obstacles to expertise has highlighted, one must be discerning about when and how to do this. This is a matter for another paper.

Chapter 3

Virtuous & Worthy Moral Deference

1 Introduction

Consider the following case:

Halloween. Joel is attending an Intro to Ethics recitation. The class is discussing the ethics of Halloween costumes, which he's never thought about before. Because Joel has a history midterm right after class, he isn't paying attention. Rather, he's reciting historical facts about minstrel shows in his head. At the end of class, Joel hears his TA state that she herself thinks donning blackface for Halloween is wrong. Joel leaves class having no clue why anyone would have this view. He also has no feelings about either blackface or minstrelsy. He knows they upset the African American community, but thinks everything seems to upset someone anyway. Still, he believes his TA. Whenever she's explained her moral views in the past, he's always been able to grasp her explanations, and has always agreed with her conclusions. Joel, a white man, decides not to paint his face black and dress up as Black Panther for

Halloween. When asked why, he says, 'It would be wrong. My TA said so. Come to think of it, I don't know why. But I want to do what's right.'

Something's off about Joel. But before I delve into our intuitions about *Halloween*, here are several things to note about the case. First, Joel's engaged in the epistemic practice of *pure moral deference*. This is the practice of forming and sustaining some moral belief *p*, solely based on someone else's purely moral testimony that *p*. What makes this pure, rather than impure, moral deference is that the testimony and deference is of an entirely moral, non-empirical nature. Let's assume that the empirical facts Joel knows are sufficient for someone else with the exact same information to figure out the wrongness of blackface on Halloween on their own. This includes facts about minstrelsy, facts about race relations in America, and the fact that blackface angers the African American community. Joel already has all the non-moral information he needs; he's just incompetent at putting these pieces of the moral puzzle together. So the new belief he accepts is purely moral, and isn't formed or sustained on the basis of any empirical information he lacks.

Second, Joel's new belief plausibly meets the standards of moral knowledge, and he's *successfully* engaged in pure moral deference. The broad consensus in the literature on pure moral deference is that this is possible. Let's say Joel's TA is indeed a reliable source of moral information, and Joel justifiably believes that he himself is pretty good at discerning good moral explanations from bad ones (he just isn't great at forming true moral beliefs on his own). The fact that he's been able to follow his

TA's explanations in the past, and the fact that he has always agreed with her conclusions, provide justification for his true belief that donning blackface for Halloween is wrong. It's no accident that Joel has stumbled upon this true belief. He demonstrates a *competence* at forming moral knowledge via pure moral deference.

Now, intuitively, Joel hasn't met the standards of an ideal moral agent. He doesn't seem to believe as *well* as he can, or act as well as he can when he acts on his new belief. This is despite the fact that his new moral belief counts as knowledge, and the fact that he's competent at deferring in order to acquire moral knowledge. This is also despite the fact that he demonstrates a concern for doing the right thing, and acquires moral knowledge in order to do so. When Joel believes that donning blackface on Halloween is wrong based on his TA's say-so, and acts on that belief, he still seems to miss out on something important.

Those in the debate over pure moral deference have various things to say about cases like *Halloween*. I take the most substantive threat against the practice to be what I call the *reasons unresponsiveness observation*. Philosophers have made this observation in various forms. The broad point is that even when we're successful in using pure moral deference to acquire moral knowledge that *p*, we can still fail to respond in valuable ways to the moral reasons that make *p* true. First, agents can be unresponsive to moral reasons by failing to have *cognitive moral understanding* of the reasons why *p* is true. Joel, for instance, doesn't intellectually grasp why all the empirical facts that he knows make donning blackface for Halloween wrong.

Second, agents can be unresponsive to reasons by failing to have *fitting affect* in response to the features that ground the moral reasons why p is true. Joel lacks the outrage that's fitting for him to have in response to the ugly facts he knows about minstrelsy.

The reasons unresponsiveness observation seems to reflect badly on the practice of pure moral deference, as well as agents who defer. Here are two potential negative implications of the observation. First, it's been argued that an agent's actions can't have *moral worth* if she doesn't have cognitive moral understanding. Second, it's been argued that being cognitively and affectively unresponsive to moral reasons signals a lack of *virtue* in the agent herself. If this is true, then Joel isn't being a virtuous moral agent, and his decision not to don blackface for Halloween doesn't have moral worth.

The main aim of this paper is to vindicate pure moral deference, and agents who defer, in the face of the reasons unresponsiveness observation. In certain contexts, it's important for agents to be motivated by a concern for doing the right thing in itself. In particular, I'll argue for the following claim:

The contextual claim about moral worth: In contexts C where an agent's unresponsiveness to the moral reasons why φ -ing is right makes her incompetent at φ -ing on the basis of those moral reasons, her act of φ -ing has moral worth if (i) she's motivated to φ by a concern for doing the right thing, and (ii) she's competent at φ -ing based on this concern.

I'll also argue for the existence of the following two virtues:

The virtue of acting for rightness' sake: An agent's competence at using her concern for doing the right thing to perform right actions, in contexts where because of her unresponsiveness to moral reasons, she'd otherwise be incompetent at doing so.

The virtue of knowing for rightness' sake: An agent's competence at using her concern for doing the right thing to form moral knowledge, in contexts where because of her unresponsiveness to moral reasons, she'd otherwise be incompetent at doing so.

Pure moral deference, practiced in the right way as Joel has, is a way for agents to meet the standards of the contextual claim. This mitigates our worries about moral worth. It's also a way for agents to exercise the virtues of acting and knowing for rightness' sake. This mitigates our worries about virtue. My account won't conclude that there's nothing missing whatsoever in agents like Joel; of course there is. But he's a much better moral agent than some would have us think.

Here's how I'll proceed. In Section 2, I'll outline the reasons unresponsiveness observation, and how it creates worries for pure moral deference and agents who practices it. In Section 3, I'll argue for the contextual claim about moral worth, and show how this mitigates worries about moral worth for deference. In Section 3, I'll argue for the existence of the two virtues of acting and knowing for rightness' sake,

and show how this mitigates worries about virtue for deference. I'll wrap up by responding to some possible objections to my account.

2 The threat of reasons unresponsiveness

In this section, I'll outline the reasons unresponsiveness observation about pure moral deference, and how it threatens the status of the practice as well as agents who engage in it. Recall that pure moral deference is the practice of forming and sustaining some moral belief p , solely based on someone else's purely moral testimony that p is true. The reasons unresponsiveness observation is that even when we're successful in using pure moral deference to acquire moral knowledge that p , we can still fail to respond in valuable ways to the moral reasons that make p true.

Here's how this section will go. First, I'll explain what I mean by moral reasons (2.1). I'll then discuss how deferring agents can fail to respond to them both cognitively and affectively (2.2). Next, I'll discuss what successful deferrers are typically responding to instead, and why (2.3). This is the moral truth, out of a non-instrumental concern for doing the right thing. Last, I'll discuss two threats that the reasons unresponsiveness observation poses (2.4). These are lack of morally worthy action, and lack of virtue.

2.1 Moral reasons

Moral reasons are what we hope to express when justifying our moral beliefs and

actions to other. They are what ground the truth of moral facts. Consider the following example. Suppose I correctly believe kicking people on a whim is wrong, and correspondingly refrain from doing so. Some statements I could use to justify my belief and actions are 'It would cause others physical pain', 'It would fail to maximise the moral good', 'Kicking people on a whim is cruel', or 'We ought to maximise the moral good'. These are all sensible things to say to someone who asks me to justify myself. These statements express moral reasons.

I'll follow the broad philosophical consensus in holding that moral reasons are facts (see e.g. Raz 1975; Scanlon 1998). The moral reasons I express above are 'the fact that kicking others on a whim would cause them physical pain', 'the fact that doing so would fail to maximise the moral good', and so on.

Moral reasons come in different levels of abstraction. The most ground-level sorts are those like 'the fact that kicking people on a whim would cause others physical pain'. These are the particular empirical facts which ground the moral status of actions or status of affairs. Those in debates about morally worthy action commonly refer to these as 'right-making reasons' or 'right-making features', 'wrong-making reasons' or 'wrong-making features', or 'morally relevant features'.⁴²

These ground-level moral reasons can ground more abstract moral reasons which are normative, as opposed to empirical, facts. For example, the empirical fact

⁴² McNaughton & Rawling (2018) refer to these as 'tier one facts', on a two-tier view of normative reasons (174). Yetter Chappell (2018) refers to these as the facts involved in 'ground-level moral explanations', as opposed to those involved in 'criterial moral explanations' (9).

that I'd cause others physical pain by kicking them is what partially grounds the normative fact that doing so would fail to maximise the moral good. The combination of empirical details about kicking others on a whim is also what grounds the fact that doing so would be cruel⁴³.

On an even more abstract level, there are the moral reasons which are normative facts about more general moral truths, untethered from the case at hand. These are moral reasons like 'the fact that we ought to maximise the moral good', or 'the fact that something is cruel is a moral reason not to do it'. These kinds of moral reasons are the ones which give the ones above their very status as moral reasons.

2.2 The reasons unresponsiveness observation

Agents who successfully practice pure moral deference to acquire moral knowledge that p can fail to respond, in various important ways, to the moral reasons that make p true. I'll focus on two such ways: cognitive moral understanding, and fitting affect.

(i) Missing out on cognitive moral understanding

The type of cognitive moral understanding I'll discuss is the type invoked by Hills (2009). Cognitive moral understanding involves a grasp of the relationship between a moral proposition p , and the reasons why p is true (101). It is a sort of know-how, rather than knowing-that. It isn't equivalent to merely knowing the facts that constitute the reasons why p is true, as Joel does in *Halloween*. Neither is it equivalent

⁴³ These are what McNaughton & Rawling (2018) refer to 'tier two facts' (174).

to believing, or even knowing, that these facts constitute the moral reasons that make p true (104).

Full cognitive moral understanding involves a grasp of all the kinds of reasons I've outlined in the previous subsection⁴⁴. However, it's implausible to think that an agent has cognitive moral understanding only if she has a maximal grasp of all such reasons. As Hills acknowledges, there's a sufficient threshold of the possession and grasp of moral reasons, above which an agent counts as having cognitive moral understanding⁴⁵.

The causal explanation for why pure moral deference fails to guarantee cognitive moral understanding is twofold. First, an agent can defer to statements like ' p ', and ' p is true', which don't contain any moral reasons as their content. Second, an agent can defer to statements which do contain moral reasons, such as ' p , and p is true because q '. But so long as she doesn't grasp the relationship between p and q , the agent still lacks cognitive moral understanding.⁴⁶

⁴⁴ Hills invokes all these kinds of reasons when she describes agents with moral understanding. (i) Ground-level moral reasons, constituted by empirical facts: The fact that animals suffer under modern farming methods is a moral reason not to eat animals (100). (ii) More abstract moral reasons, constituted by normative facts grounded by empirical facts: The fact that someone is a person is a moral reason not to kill her (115). (iii) Moral reasons constituted by facts about general moral obligations: The question of whether we ought to maximise welfare is relevant to understanding why eating meat is wrong (100).

⁴⁵ For Hills, although moral understanding requires 'at least to some extent, a systematic grasp of morality', and that 'there might be a cutoff point before which you do not count as having moral understanding', this is not to say that 'you need to have a grasp of anything that could be morally important' (Hills 2009, 103).

⁴⁶ Additionally, the problem isn't just that pure moral deference 'doesn't guarantee' cognitively moral understanding. Pure moral deference seems to be mutually exclusive with grasping the relationship between p and the moral reasons why p is true. Cognitive moral understanding is presumably the sort of thing for which once you have it, your belief typically comes to be substantively based on this understanding.

(ii) *Missing out on fitting affect*

The second way in which a successful deferring agent can fail to respond to moral reasons is by failing to demonstrate fitting affect in response to the properties that generate these moral reasons. For example, one moral reason why donning blackface on Halloween is wrong is the fact that it causes the African American community to suffer. When one witnesses this suffering, or imagines it, it's fitting to respond with emotions like outrage on the community's behalf. Most of us think there's something sub-par about someone who doesn't, like Joel.

Typically, we experience moral affect in response to the ground-level properties of cases, such as particular instances of suffering. What about features which generate more abstract reasons, like cruelty? It's controversial as to whether we typically do, or can, directly respond to such things with affect.⁴⁷ In any case, nothing much about what I'll say hinges on this possibility.

Pure moral deference, obviously, fails to guarantee emotional responses of this nature (e.g. Howell 2014; Enoch 2014; Lord FC). There's something inappropriate, even, about an agent who acquires affective reactions immediately upon accepting someone else's moral testimony.

Here's one last note. On Hills' view, moral understanding is cognitive. But

⁴⁷ My own view is that one can directly respond emotionally to higher-level moral features like cruelty. Not everyone will agree that this is possible, or agree that this is typical. For instance, Enoch (2014) notes that emotionally responding to 'second-order wrongs' like 'the subjection of others to unacceptably high risk' is atypical, although possible (28).

other views hold that moral understanding involves precisely the type of affective responses I'm discussing here (e.g. Lord FC). So, deference doesn't just fail to guarantee the cognitive type of moral understanding Hills adopts. It also fails to guarantee moral understanding on any view which holds that affective responses to the properties which generate moral reasons are necessary.

2.3 Responding to moral truth, out of a concern for doing the right thing

We've seen how agents can successfully practice pure moral deference, but still fail to respond to moral reasons in important ways. But agents who are competent at forming knowledge via deference must be responding to *something*. What is this something, and why are they responding to it?

Perhaps successful deferrers are responding to moral testimony itself. One thing we can say Joel is responding to is 'his TA's say-so' when he defers. But merely responding to moral testimony isn't good enough for deference to transmit moral knowledge. Not just any old testimony, picked out in any old way, will do. In order for deference to be successful, agents must be aiming, and be competent and successful at aiming, at some further feature that moral testimony can have.

What successful deferrers are aiming at is the moral truth. And agents who aim for the moral truth typically do so out of a non-instrumental *concern for doing the right thing* in itself. Although deference is an epistemic practice, beliefs guide action, and one is typically interested in acquiring true moral beliefs so that one can perform right actions. I'll follow Sliwa (2016) in adopting the following definition of what a

'concern for doing the right thing' is. To demonstrate 'concern' is to experience conative states such as desires and intentions with respect to what one is concerned about. To demonstrate concern for doing what's right includes having desires and intentions etc. to do what's right. This concern is non-instrumental; an agent who is concerned about doing what's right cares about doing what's right for its own sake. As for 'doing the right thing', this includes not just performing actions that are morally required, but also those that are supererogatory and permissible (395-396).

I say that truth-aiming moral agents are 'typically' demonstrating a concern for doing the right thing, because it's also possible to demonstrate an independent non-instrumental concern for believing moral truths (or truths in general). I think it's quite common for agents to simultaneously be concerned with both doing the right thing in itself, and believing the truth in itself. But I lack the space to discuss this here, and will set the matter aside.

2.4 The threat of the reasons unresponsiveness observation

At this point, readers might be persuaded that failing to have cognitive moral understanding and fitting affect is bad in itself. But the case against pure moral deference, and agents who defer, becomes even stronger when we consider some more negative implications that seem to follow. Cognitive moral understanding and fitting affect seem to ground other valuable moral goods. It's been argued that such goods fail to be demonstrated or acquired when an agent succeeds at pure moral deference. I'll focus on two such candidates: morally worthy action, and virtue.

(i) *Lack of morally worthy action*

Moral worth is a property of actions. What it means for an action to have moral worth, roughly, is that the agent deserves credit for her success in performing that right action (Sliwa 2016, 395). An action being right doesn't suffice for moral worth; the action must have been performed *well*. An agent who performs some right action by accident, for instance, doesn't perform actions with moral worth.

The worry for pure moral deference is that when agents act on deferential beliefs, those actions fail to have moral worth, because they fail to be performed in response to moral reasons that make that action right. According to Hills, actions are morally worthy only if they are right actions performed for the right reasons, which one has cognitively grasped as reasons to act (113-119).

(ii) *Lack of virtue*

Virtue is a property of agents. While philosophers disagree on what virtue requires, most agree on at least two things. First, an agent counts as maximally virtuous only if she has some specified kind, or kinds, of sensitivity towards moral reasons. Second, this specified kind, or kinds, of sensitivity is either of a cognitive sort, an affective sort, or both. If both of these claims are true, then the reasons unresponsiveness observation threatens to mean that successful deferrers aren't virtuous.

Hills is explicit that cognitive moral understanding is necessary for virtue, and that this is something that makes the practice of pure moral deference bad (108-113). But she also indicates that other types of responses to moral reasons are required for virtue, including affect (112).⁴⁸ Howell is similar to Hills, in the sense that he thinks full virtue involves (amongst other things) both cognitive moral understanding (406-407), and fitting affect towards e.g. suffering (404-405). But unlike Hills, he argues that it's the deferring agent herself that's sub-optimal, rather than the practice of deference itself (409).

3 The contextual claim about moral worth

So far, I've discussed the reasons unresponsiveness observation for pure moral deference, and how this threatens to mean that successful deferring agents can't perform morally worthy actions, and lack virtue.

I'll now begin my defense of pure moral deference and agents who practice it. In this section, I'll show how it's possible for successful deferring agents to act on their deferential knowledge with moral worth. This will involve vindicating the practice of being motivated to act by a concern for doing the right thing, as opposed to being motivated by moral reasons, in certain contexts. In particular, I'll argue for

⁴⁸ According to Hills, 'a good, virtuous person is someone whose whole self—her thoughts, decisions, feelings, and emotions as well as her actions—is structured by her sensitivity to [moral reasons]' (Hills 2009, 112).

the following claim:

The contextual claim about moral worth: In contexts C where an agent's unresponsiveness to the moral reasons why φ -ing is right makes her incompetent at φ -ing on the basis of those moral reasons, her act of φ -ing has moral worth if (i) she's motivated to φ by a concern for doing the right thing, and (ii) she's competent at φ -ing based on this concern.

Here's a note about this claim. Moral worth comes in degrees. All that the contextual claim says is that an agent who fulfils these requirements performs an action with *some* moral worth. My goal isn't to argue such an action has *maximal* moral worth.

For now, I'll start by outlining the existing debate over whether morally worthy actions must be motivated by a concern for doing the right thing, or a concern for moral reasons (3.1). Next, I'll argue for part (i) of the contextual claim (3.2). Then, I'll argue for part (ii) of the contextual claim (3.3). Next, I'll explain why I define contexts C in terms of incompetence (3.4).

Finally, I'll explain how this mitigates worries about moral worth for pure moral deference, in the face of the reasons unresponsiveness observation (3.5). Successfully practicing pure moral deference, and acting on the basis of deferential knowledge, is one way of meeting the requirements of the contextual claim. So, it is one way of performing actions with moral worth.

3.1 The motives for action debate

Quite separately of the debate on pure moral deference, there's another lively philosophical debate over what sort of motives are necessary or sufficient for an agent's actions to have moral worth.

This debate has its origins in Kant (1785), who argued that an agent acts with moral worth only if she's motivated to act out of a concern for doing the right thing, or by the very fact that the action is right. This is the type of motivation that Joel demonstrates when he listens to his TA, and acts based on her testimony. In contrast, Smith's (1994) well-known argument against externalist accounts of reasons for action held that good agents act directly on the basis of right-making features, rather than a non-instrumental concern for doing what's right.

Here are two necessary claims one can make about motivations and moral worth. Philosophers hold these claims in various combinations:

- (1) In all contexts, an agent's right act of φ -ing has moral worth only if her φ -ing is motivated by the moral reasons that make φ -ing right.
- (2) In all contexts, an agent's right act of φ -ing has moral worth only if her φ -ing is motivated by a concern for doing the right thing.

One could accept (1), but reject (2). In other words, morally worthy actions are performed on the basis of moral reasons, and a concern for doing the right thing isn't necessary (e.g. Smith 1994, Arpaly 2002 & 2003, Markovits 2010). Alternatively, one

could accept (2), and reject (1). On this view, morally worthy actions are motivated by a concern for doing the right thing, and acting on the basis of moral reasons isn't necessary (e.g. Sliwa 2016, Johnson King FC). Last but not least, the most demanding view accepts both (1) and (2). On this view, morally worthy actions are motivated both by a concern for doing the right thing, and by moral reasons (Singh 2018).

I reject claim (1), and remain neutral on (2). What I'll argue for in this section is a *contextual* claim about what sort of motivation is sufficient in contexts which fit a certain description. Here's the claim I make:

The contextual claim about moral worth: In contexts C where an agent's unresponsiveness to the moral reasons why φ -ing is right makes her incompetent at φ -ing on the basis of those moral reasons, her act of φ -ing has moral worth if (i) she is motivated to φ by a concern for doing the right thing, and (ii) she is competent at φ -ing based on this concern.

I'll now proceed to defend this contextual claim.

3.2 Defending the contextual claim: part (i)

I'll start by motivating part (i) of the claim with some cases. Here's a case adapted from Lillehammer (1997):

Temptation. Teresa is tired of her husband. At a party, she meets a charming woman and is tempted to have an affair. She judges that this would be wrong on account of her husband's feelings. She's temporarily indifferent to his

feelings, however, because she's still mad with him for not helping out with the dishes last night. Out of a concern for doing the right thing, Tammy doesn't have the affair. (Adapted from Lillehammer 1997, 192)

Teresa's husband's feelings generate a moral reason for her not to have the affair. She is conatively and affectively unresponsive to this reason. This unresponsiveness makes her unmotivated, and hence incompetent, at doing the right thing based on that reason. This is despite the fact that Teresa has cognitive moral understanding of why her husband's feelings makes having the affair wrong. One can be responsive to moral reasons in some ways, but unresponsive to them in other ways, that overall decreases one's motivation to do the right thing based on those reasons.

Teresa's concern for doing the right thing stops her from doing the wrong thing. Intuitively, her action has some amount of moral worth. This is despite the fact that she experiences recalcitrant emotions and desires. We might even think that her action has moral worth *because* of the fact that she experiences recalcitrant emotions and desires, and demonstrates the admirable willpower not to act on them. This lends intuitive support to the contextual claim. Consider the next case, which does the same:

Changing One's Mind. Before today, Chuck has always believed that morality isn't very demanding in terms of individual sacrifice. But after reading *Famine, Affluence, and Morality*, he can't help concluding that he's morally required to donate much of what he has to the distant needy. He's still

reluctant to do this. And despite trying his best to imagine the suffering of the distant needy, his imagination isn't good enough to generate much emotion. Despite this, and out of concern for doing the right thing, he proceeds to donate half of his savings to charity. (Adapted from Lillehammer 1997, 191)

Like Teresa, Chuck's action intuitively has some amount of moral worth. Many agents in contexts *C* will be like Teresa and Chuck. They will have cognitive moral understanding of moral reasons, but will lack the motivation to perform right actions because they lack the appropriate emotions and desires. This is no surprise, given the links between emotions, desires, and motivation. Not all cases are like this, however:

Humility. Huifen loves the film *Crazy Rich Asians*, and loves the fact that it's a milestone for Asian representation in Western media. But many of her Indian friends have told her that the way the film depicts South Asians is a problem. Huifen has noticed that the only South Asian characters in the movie are servants, but doesn't see any problem with this. But she knows that light-skinned Chinese people like her tend to have moral blind spots when it comes to racism and colourism in the Asian community, and have interests in believing there's nothing problematic about the film. Suspicious of herself, Huifen suspends her immediate beliefs on the matter, and has the higher-order belief that her friends are more likely to be correct. Out of concern for doing the right thing, she shares an article on Facebook, written by one of her

friends, highlighting the racist aspects of the film.

Unlike Teresa and Chuck, Huifen lacks cognitive moral understanding of the reasons why there's anything morally problematic with *Crazy Rich Asians*. Because of her unresponsiveness to moral reasons when it comes to racism and colourism in the Asian community, she is incompetent at doing the right thing based on those reasons. But because she recognises this, she uses her concern for doing the right thing to perform the right action. Intuitively, her act of spreading awareness of the issue on social media has some amount of moral worth.

3.3 Defending the contextual claim: part (ii)

According to the contextual claim about moral worth, it isn't enough for reasons-unresponsive agents to be motivated by a concern for doing the right thing. For her right act of φ -ing to have moral worth, she must also fulfil part (ii) of the claim. She must be *competent* at φ -ing based on her concern for doing the right thing.

This is because agents who who are incompetent at acting for moral reasons, but are motivated by doing a concern for doing the right thing, can accidentally perform right actions. If not for part (ii) of the contextual claim, the claim would incorrectly predict that these actions have moral worth. Consider the following case:

Reckless Uncertainty. Umar is at the ballot box deciding whether to vote to abolish the death penalty. Although he wants to do the right thing, he can't figure out what that is. A flash of divine intuition tells him that flipping a coin

will reveal the right course of action. Based on the result of a coin toss, Umar votes to abolish the death penalty.

Although Umar does the right thing, his action isn't morally worthy. He does the right thing by accident, as a result of a reckless decision-making procedure. He believes this is a good procedure, but it isn't. Umar is incompetent at acting based on his concern for doing the right thing. Without part (ii) of the contextual claim, his action would incorrectly count as morally worthy, since he was motivated to act by a concern for doing the right thing, and performed the right action.

3.4 Defending the contextual claim: contexts C

Here, I'll explain why I define the scope of the contextual claim as I do. Recall that parts (i) and (ii) of the contextual claim only apply in contexts C , which are contexts where an agent's unresponsiveness to the moral reasons why φ -ing is right makes her incompetent at φ -ing on the basis of those moral reasons.

Why define C in terms of incompetence? Why not define C counterfactually, and say that the relevant contexts are those in which an agent's unresponsiveness to moral reasons 'would otherwise cause her to perform the wrong action'?

Again, the answer has to do with agents who accidentally perform the right actions. An agent can be incompetent at acting for moral reasons, lack any concern for doing the right thing, and nonetheless do the right thing by accident. The contextual claim should allow us to assess the actions of such agents, and predict

that their actions lack moral worth. Defining *C* counterfactually doesn't allow us to do this, and defining *C* in terms of incompetence does. Here's an example:

Unsatisfied Temptation. Tammy is tired of her partner. At a party, she meets Teresa, and is tempted to have an affair. Tammy judges that this would be wrong on account of her partner's feelings, but she's temporarily indifferent to them. Sparks are flying, and they can both tell that the other party wants the affair. Tammy is about to give in to her temptation. But before she makes an explicit move, Teresa says, 'Let's not do this. I have a husband, and it would be wrong.' Tammy rolls her eyes and walks away, deciding not to pursue the matter any further.

Tammy shares important similarities with Teresa, from the original *Temptation* case. Both of them are unmotivated by moral reasons. Both are also in a position to perform a morally worthy action, by demonstrating a concern for doing the right thing, and competently acting on this concern. It's just that Teresa does, and Tammy doesn't. When we judge the moral worth of Tammy's actions, we should hold them up to the same standards as Teresa's actions—that is, to the standards of parts (i) and (ii) of the contextual claim.

Defining *C* counterfactually renders the contextual claim silent on agents like Tammy. Her unresponsiveness to moral reasons would not have caused her to perform the wrong action, since she doesn't in fact perform the wrong action.

In contrast, defining *C* as contexts in which agents are incompetent at acting

on the basis of moral reasons delivers not just a verdict, but the correct verdict that Tammy's actions lack moral worth. Tammy's unresponsiveness to moral reasons makes her incompetent at performing right actions on the basis of them. If not for Teresa's reluctance, Tammy would have initiated the affair of a disregard for her partner's feelings. Her doing the right thing is a lucky accident. The contextual claim, where C is defined in terms of incompetence, correctly predicts that Tammy's right action has no moral worth.

3.5 Morally worthy action & pure moral deference

In the subsections above, I've argued for each component of the following claim:

The contextual claim about moral worth: In contexts C where an agent's unresponsiveness to the moral reasons why φ -ing is right makes her incompetent at φ -ing on the basis of those moral reasons, her act of φ -ing has moral worth if (i) she's motivated to φ by a concern for doing the right thing, and (ii) she's competent at φ -ing based on this concern.

Now, I'll explain how the contextual claim makes it possible for agents who successfully practice pure moral deference to act on their deferential knowledge with moral worth.

Recall Joel from our original *Halloween* case. Notice that he and his actions fulfil the requirements of the contextual claim. First, Joel's situation falls under the category of contexts C . He's unresponsive to the moral reasons why donning

blackface for Halloween is wrong, and is incompetent at refraining from doing so on the basis of those reasons. It's only because of his TA's moral testimony that he makes the decision to do this. Second, Joel fulfils requirement (i) of the contextual claim, by being motivated to refrain from donning blackface for Halloween by a concern for doing the right thing.

Third, Joel fulfils requirement (ii) of the contextual claim: he is competent at doing the right thing based on this concern. By successfully practicing pure moral deference, he has demonstrated a competence at acquiring moral knowledge via deference. His deferential moral knowledge is what spurs him to do the right thing. So, his competence at acquiring moral knowledge via deference gives him the competence to perform the right action based on a concern for doing the right thing, when he acts on his deferential knowledge.

So, successfully practicing pure moral deference, and acting on deferential knowledge, is one way of meeting all the requirements of the contextual claim about moral worth. Just as Tammy's, Chuck's, and Huifen's actions have moral worth, so too does Joel's choice not to don blackface for Halloween. This mitigates worries about morally worthy action in the face of the reasons unresponsiveness observation. When an agent who successfully defers φ -s on the basis of her deferential moral knowledge, her act of φ -ing can have some amount of moral worth.

4 Virtues of acting & knowing for rightness' sake

In the previous section, I argued for the contextual claim about moral worth. I also showed how the contextual claim mitigates worries about morally worthy action for pure moral deference, in the face of the reasons unresponsiveness observation.

Now, I'll argue that the contextual claim implies that the following two moral virtues exist:

The virtue of acting for rightness' sake: An agent's competence at using her concern for doing the right thing to perform right actions, in contexts where because of her unresponsiveness to moral reasons, she'd otherwise be incompetent at doing so.

The virtue of knowing for rightness' sake: An agent's competence at using her concern for doing the right thing to form moral knowledge, in contexts where because of her unresponsiveness to moral reasons, she'd otherwise be incompetent at doing so.

Here's a preliminary note about these virtues, similar to the one I made about the contextual claim. My goal isn't to argue that agents who exercise these two virtues are being *maximally* virtuous. Some virtues are more ideal than others, in the sense that they are virtues which would be possessed and exercised by ideal moral agents. The virtues of acting and knowing for rightness' sake don't fit this bill, since they can

only be exercised by flawed agents. Readers might worry that this means they can't legitimately count as virtues; I'll address this worry in the conclusion.

I'll start by arguing for these virtues in turn (4.1, 4.2). Then, I'll explain how they mitigate worries about virtue for pure moral deference in the face of the reasons unresponsiveness observation, and also allows the epistemic act of deferring itself to have moral worth (4.3).

4.1 The virtue of acting for rightness' sake

We've seen how cases of moral uncertainty are opportunities for agents to exercise their concern for doing the right thing. *Reckless Uncertainty* was one such (failed) case. Sepielli (2016) notes what morally uncertain agents should do:

'When you're confident enough about which considerations would render your actions objectively right, you act directly on those considerations; but when you're not confident, you ... aim for objective value itself. ... When I am consciously uncertain about morality, it seems that I must either act on [a responsible] uncertaintist norm, or else do an action that is not fully norm-guided ... [like] a "leap of faith". But, just intuitively, the motivations of an agent in taking a leap of faith can hardly be called *virtuous*; such an agent seems to be acting recklessly by her own lights. (Sepielli 2016, 11-12, emphasis mine)⁴⁹

⁴⁹ The particular 'uncertaintist norm' that Sepielli's addressing is the practice of moral hedging. I won't go into that debate here. But I'll point out that another wise uncertaintist decision-making

Sepielli's quote here is instructive, in that it links *virtue* to the debate about acting for moral reasons versus acting out of a concern for doing the right thing (acting on 'which considerations would render your actions objectively right' versus 'aiming for objective value itself'). Sepielli (perhaps loosely) speaks of virtue as being a feature of an agent's motives. In contrast, I think that virtues properly reside in the agent herself. More specifically, a virtue is a moral competence an agent possesses, which when exercised, confers moral worth upon her actions.

As we've seen in the previous section, this is precisely the type of competence possessed and exercised by agents like Teresa from *Temptation*, Chuck from *Changing One's Mind*, and Huifen from *Humility*. The same goes for Joel, in the *Halloween* case we started with. By fulfilling the standards of the contextual claim about moral worth, they've signalled the following virtue:

The virtue of acting for rightness' sake: An agent's competence at using her concern for doing the right thing to perform right actions, in contexts where because of her unresponsiveness to moral reasons, she'd otherwise be incompetent at doing so.

4.2 The virtue of knowing for rightness' sake

In addition to the virtue of *acting* for rightness' sake, we should also think a parallel moral virtue exists for moral *knowledge*:

procedure is the practice of asking someone else what to do, where you have good reason to think that this she'll make the right call.

The virtue of knowing for rightness' sake: An agent's competence at using her concern for doing the right thing to form moral knowledge, in contexts where because of her unresponsiveness to moral reasons, she'd otherwise be incompetent at doing so.

Unlike the virtue of acting for rightness' sake, this moral 'knowing-virtue' isn't exercised by Teresa and Chuck. They already know what the right thing is to do, based on their cognitive grasp of moral reasons. Huifen doesn't exercise this virtue either, since she chooses to suspend judgment rather than believe her friends' testimony. Joel is the only agent who exercises this virtue.

Here's why we should think the virtue of knowing for rightness' sake exists. Just as good moral agents are competent at performing right actions, they're also competent at acquiring moral knowledge. One might think that ideally, this competence should be manifested in the form of responsiveness to moral reasons. Yet, just as it is in the domain of action, there are many contexts in which we are epistemically unresponsive to moral reasons. This is true of Huifen, Umar, and Joel. It's important for moral agents to be good at acquiring moral knowledge in such situations.

Importantly, the virtue of knowing for rightness' sake is a *moral* virtue, and not just an epistemic one. This is because moral knowledge, in addition to having epistemic value, has moral value as well. One possibility is that moral knowledge is

intrinsically morally valuable.⁵⁰ I find this appealing, but it complicates the picture and I won't discuss it here.⁵¹

Setting this aside, moral knowledge is certainly *instrumental* to a whole host of other things which are themselves morally valuable. Having moral knowledge is a useful epistemic attitude to adopt, as it is conducive to the following outcomes. (1) First, having moral knowledge is valuable because it allows us to perform right actions in a competent, morally worthy way. We've seen how Joel has used his moral knowledge to competently refrain from donning blackface on Halloween. (2) Second, having moral knowledge can help us become responsive to moral reasons in the long run. There are many situations in which an agent, had she not first accepted some true moral belief *p*, wouldn't later have achieved cognitive moral understanding of the moral reasons why *p*, fitting emotional responses to the features generating moral reasons why *p*, fitting conative and motivational responses towards these features, and so on. Cases of moral learning in children demonstrate this the most clearly, but it's easy to overlook that this frequently happens to adult agents too. (3) Third, insofar as these responses to moral reasons are important for other kinds of virtues, having moral knowledge can help us develop these virtues in

⁵⁰ Suppose an agent is in a position to acquire a piece of moral knowledge, without being in a position to acquire any other important good like cognitive moral understanding, affective reactions, and so on. Such a view would hold that even if this agent would never be in a position to act on this piece of moral knowledge, she would still be better off having it than not.

⁵¹ If moral knowledge is intrinsically valuable, this indicates the existence of quite a separate virtue we might call virtue of 'knowing for moral truths' sake': An agent's competence at using her concern for *having moral knowledge* to form moral knowledge, in contexts where because of her unresponsiveness to moral reasons, she'd otherwise be incompetent at doing so. Or, the value of both the action-virtue and the knowing-virtue might really derive from a more fundamental value of something like (roughly) 'morality itself'.

the long run too. (4) Fourth, insofar as these responses to moral reasons might help us perform actions with more moral worth, having moral knowledge can help us to achieve this too in the long run.

To be clear, moral knowledge isn't the only epistemic attitude which is conducive to the outcomes above. Consider Huifen. She lacks moral knowledge, but adopts the epistemic attitude of believing that her friends are more likely to be correct. This allows her to perform the right action in a competent, morally worthy way. Realising that her friends are more likely to be correct might also spur her to adopt an interrogative type of epistemic attitude, where she is motivated to understand the reasons why there might be a problem with Crazy Rich Asians. It's enough, however, that having moral knowledge can play the same role as having higher-order epistemic attitudes about what's likely to be correct. So long as having moral knowledge can play this causal role, it has moral value.

There are, however, some things that moral knowledge brings to the table that other epistemic attitudes cannot. (5) For one, having moral knowledge allows agents to competently perform *sincere* right actions. Epistemic attitudes which aren't true, justified, first-order moral beliefs can't do this. Even though Huifen has done something good, her act of sharing her friend's article on Facebook still seems weird—it isn't completely sincere. It seems *better* for an action to be performed on the basis of an agent's actual beliefs, rather than, say, her higher-order assessment that the action is likely to be right. It's appropriate for Huifen's friends to feel slighted. And

this isn't just because they know she has good reason to believe their testimony, instead of suspending judgment. Consider the following case:

Monster in the Mines. Harvey has some serious childhood trauma. While wandering into his father's mines as a young boy, he saw a terrifying monster. Harvey wants to return to the mines to see if the monster is still there. He tells his girlfriend Sabrina about his plans. Unbeknownst to Harvey, Sabrina has a good reason to think he's right about what he saw: she's a witch who's seen many supernatural creatures herself. Because revealing her secret to Harvey would put him in danger, Sabrina pretends she doesn't believe him. She raises her eyebrows and says, 'Well, you could be wrong, or you could be right. But I fully support you. Let's go into the mines together to confront what you saw... or what you *think* you saw.' Although Harvey feels grateful, he's mostly just annoyed. He says, 'Thanks, but no thanks. I don't need your help.'

Although it's fitting for Harvey to feel grateful, it's also fitting for him to feel slighted. This is despite the fact that he doesn't know that Sabrina has a good reason to believe him, and despite the fact that he knows that what he's saying sounds crazy. From Harvey's perspective, Sabrina's offer to accompany him to confront his trauma isn't sincere, since she appears not to believe that he really saw a monster in the mines. Her claim that she fully supports him doesn't appear to be sincere either. Sabrina is being patronising to Harvey. From his perspective, it would be better for her to accompany him into the mines out of a full-blooded, sincere belief that he'd

really seen a monster.

Agents like Huifen and Sabrina demonstrate how moral knowledge brings something to the table that other epistemic attitudes can't. This is because merely performing right actions without knowing, or at least believing, that they're true doesn't allow us to be sincere. The value of sincerity might be cashed out in various ways. One plausible way is to hold, as many of us do, that sincerity itself is a virtue.⁵² If so, then Huifen's act of sharing her friend's article would be *more* morally worthy, and would have signalled more virtue, had it stemmed from moral knowledge rather than her higher-order probabilistic assessment.

4.3 Virtuous pure moral deference

To recap, I've argued for these two virtues in the subsections above:

The virtue of acting for rightness' sake: An agent's competence at using her concern for doing the right thing to perform right actions, in contexts where because of her unresponsiveness to moral reasons, she'd otherwise be incompetent at doing so.

The virtue of knowing for rightness' sake: An agent's competence at using her concern for doing the right thing to form moral knowledge, in contexts where because of her unresponsiveness to moral reasons, she'd otherwise be

⁵² Or, more precisely, that an agent's competence at acting (and speaking, communicating, etc.) in accordance with her actual beliefs is a virtue.

incompetent at doing so.

If these virtues exist, it's easy to see how agents who successfully practice pure moral deference, and act on their deferential moral knowledge, are signalling virtue. This mitigates worries about virtue for pure moral deference in the face of the reasons unresponsiveness observation.

The first reason is that acting on deferential moral knowledge is a way of exercising the virtue of *acting for rightness' sake*. Joel has done this, by deferring to his TA's testimony that donning blackface for Halloween is wrong, and then acting on his newfound moral knowledge. The second reason is that the very epistemic act of successfully practicing pure moral deference is a way of exercising the virtue of *knowing for rightness' sake*. Joel has done this with his very act of deferring to his TA's testimony, in the competent way that he has. When Joel defers, and acts on his deferential knowledge, he is signalling virtue after all.

The virtue of knowing for rightness' sake also helps to mitigate the separate worry about moral worth in the face of the reasons unresponsiveness observation. In Section 3, I argued that acting on deferential moral knowledge can be morally worthy. But the virtue of knowing for rightness' sake means that the very act of epistemically deferring can be morally worthy too. This is because actions produced by exercising virtues have moral worth. When an agent successfully practices pure moral deference, and exercises the virtue of knowing for rightness' sake, her act of deferring has moral worth.

5 Conclusion

The main goal of this paper was to vindicate pure moral deference, and agents who defer, in the face of the reasons unresponsiveness observation. I've done this in Sections 3 and 4, by showing how agents who successfully defer can perform actions with moral worth, and signal virtue.

Let's return to the case of *Halloween* we started with. What I've argued for delivers the following verdicts. When Joel defers to his TA's say-so and acquires moral knowledge that donning blackface for Halloween is wrong, he exercises the virtue of knowing for rightness' sake, and his act of deference itself is morally worthy. When Joel acts on his deferential belief by refraining to don blackface for Halloween, he exercises the virtue of acting for rightness' sake, and this deferential action has moral worth too. Joel isn't perfect: he isn't acting with full moral worth, and he isn't exhibiting ideal virtues and being maximally virtuous. Still, it isn't as if he's acting with *no* moral worth and signalling no virtue at all.

I'll wrap up by responding to some potential objections to my account.

(1) *Objection:* The virtues of knowing and acting for rightness' sake can't be real virtues, since they presuppose an existing flaw in the agent.

Response: The first thing to note is that this presupposes a view of virtue that's

far too demanding. Perhaps the ideal moral agent is always responsive to moral reasons, in all possible valuable ways, with respect to all possible moral questions. But why peg our standards for virtue to such an agent? As a matter of fact, it's impossible for human agents to achieve this ideal, owing to our cognitive, emotional, and even neurobiological limitations. Moreover, as I argue in the chapters before this, the limited nature of our experiences place severe limits on our moral epistemic capacities, often through no fault of our own. In light of this, our standards for virtue cannot possibly be that high.⁵³ And the idea that our standards for virtue can't be pegged to absolute perfection isn't new. The ancient Stoics, for instance, held that the standard of Sage, all of whose actions are fully virtuous, is the highest standard achievable by finite, imperfect, and limited agents.

Moreover, the fact that virtues can presuppose existing flaws in agents is, really, a positively good thing. Our flaws don't render us utterly helpless: we can adopt wise, morally valuable strategies that help us do the right things in spite of them. This includes having a standing concern for doing the right thing in itself, and being competent at gaining moral knowledge via social practices like pure moral deference. By acknowledging the essentially non-ideal nature of human moral agents, my account gives us space to strive towards becoming *good* non-ideal moral agents.

⁵³ The idea that our standards for virtue can't be pegged to absolute perfection isn't new. The Stoics, for instance, held that the standard of the Sage, all of whose actions are fully virtuous, is the 'highest standard achievable by *finite, imperfect, and limited*' agents (Vazquez MS, emphasis mine).

(2) *Objection:* My account makes virtue and moral worth too cheap.

Response: The virtues of knowing and acting for rightness' sake are actually very demanding. The moral worth that comes with exercising these virtues doesn't come easily. Being competent at acquiring moral knowledge via testimony involves meeting substantive moral and epistemic demands.

First, agents must be competent at identifying moral experts. Consider Joel. One reason why he justifiably trusts his TA is that he's been able to grasp her explanations in the past. This presupposes that he's already competent, to a substantial degree, at responding to moral reasons. Sliwa (2017) notes that without some moral competence of our own, we can't judge the moral competence of others (551). Second, one needs to be able to identify moral experts who are trustworthy. Sliwa notes again that one needs to be sensitive to traits like sincerity, integrity, a concern for helping others, and so on (551).

Third, in order to possess and exercise the virtues of knowing and acting for rightness' sake, an agent has to be competent at identifying contexts in which she's unresponsive to moral reasons, and thus incompetent at performing right actions on the basis of those reasons. This involves knowing what one's moral blind spots are. This is easy when you're consciously uncertain about some matter: Umar, for instance, knows that he's unresponsive to moral reasons, since he knows that he can't figure things out on his own. But knowing that you have a moral blind spot is

much more difficult when you're responsive to illusory moral reasons, responsive to moral reasons in the wrong way, or lack knowledge altogether that there are moral reasons you aren't responding to. Consider Huifen. She knows that her privileged social location makes her bad at responding to moral reasons with respect to racism in the Asian community. This is not an easy realisation to come to. Moreover, many others in her position would be biased towards positively thinking that they have cognitive moral understanding of why the negative representation of South Asians in *Crazy Rich Asians* is not morally problematic. Such agents, for example, might take the fact that South Asians are indeed often servants in Singapore as a moral reason to think that there is nothing wrong with the film. It would be even harder for such agents to recognise that their convictions might be false.

Fourth, agents in these positions have to know who to defer to in order to acquire moral knowledge. Knowing who to believe is even harder when you're flying blind. Deferring to members of oppressed groups with respect to certain matters, as Huifen has done, isn't an immediately obvious strategy to many people in socially privileged positions.

The two objections above were ones levelled at specific details of my account. Now, I'll respond to two objections are ones levelled at optimistic accounts of pure moral deference in general.

(3) *Objection:* Optimistic accounts of pure moral deference, like mine, license epistemic laziness. If we're permitted to defer, what would be the problem with doing so all the time?

Response: My account doesn't license epistemic laziness. Consider the following case, which is presumably one that this type of objection worries about:

Competent But Lazy. I'm wondering whether the collective action problem of factory farming makes it morally permissible for me to eat meat. I could probably figure this out on my own. I'm a competent philosophy PhD student, and emotionally sensitive to non-human animals. I have extra reason to think I have it in me, because something I wrote years ago on the topic was published in a food ethics textbook. I've just completely forgotten what my arguments and conclusions were, as I've since moved on to epistemology. I'm tired of thinking about animals. So, I ask my morally reliable, trustworthy friend Eilidh whether I'm permitted to eat meat in the face of the collective action problem. She tells me that it isn't. I justifiably accept this true belief, and now have moral knowledge. I don't think about the issue any further.

What I've done seems at least distasteful. Fortunately, my particular account doesn't condone this type of epistemic laziness. The virtues I've argued for are contextual, in the sense that they only apply when the agent is in a very bad position to acquire moral knowledge on her own. My account doesn't imply that we can outsource our moral beliefs to others when we *are* in good positions to form moral knowledge by

ourselves. If one thinks that knowing and acting in response to moral reasons is valuable, then one can agree with my account, while holding that it would be much better of me to figure things out on my own.

(4) *Objection:* Agents who practice pure moral deference will tend to end up with a bunch of bare moral knowledge, but not much else.

Response: This worry is unfounded. Nothing about my account implies that after we defer, it isn't then desirable to then attempt to achieve goods like cognitive moral understanding, fitting affect, desires, motivations, and so on. I agree that this is true, since having these goods brings us closer to being an ideal moral agent. But where I disagree with philosophers like Hills is the *order* in which it's appropriate to acquire moral knowledge and these other moral goods. Hills' rule of thumb is that moral agents should acquire cognitive moral understanding of the reasons why p , before believing and hence knowing that p . But in many contexts, I think it's perfectly acceptable for us to do things the other way around. We can justifiably accept a belief via testimony in order to form moral knowledge, and then proceed to acquire cognitive moral understanding, and other moral goods like affect, desires, and motivations.

Bibliography

- Antony, L. (1993). Quine As Feminist: The Radical Import of Naturalized Epistemology. In Antony, L. & Witt, C. (Eds.), *A Mind of One's Own: Feminist Essays on Reason and Objectivity*, 185-226. Boulder: Westview Press.
- Arpaly, N. (2002). Moral Worth. *Journal of Philosophy*, 99(5), 223-254.
- Arpaly, N. (2003). *Unprincipled Virtue*. New York, NY: Oxford University Press.
- Arpaly, N. (2018). *Deliberation and Fetish*. Workshop presentation, Washington University Workshop in Ethics, St. Louis, Missouri, USA, October 12 2018.
- Barnes, B. (Jan 23, 2015). 'The Wolfpack' Tells of One New York Apartment With Seven Children Locked Inside. *The New York Times*. Retrieved from <https://nyti.ms/2Gh8cqX> on March 8 2017.
- Blinder, A. & Fernandez, M. (April 23, 2017). Bearing Witnessed to Executions: Last Breaths and Lasting Impressions. *The New York Times*. Retrieved from <https://nyti.ms/2p8QPBs> on May 9 2017.
- Brooks, D. (Feb 1, 2018). The Abortion Memo. *The New York Times*. Retrieved from <https://nyti.ms/2FBWFAY> on Feb 15 2018.
- Campbell, R. (1998). *Illusions of Paradox*. Boulder: Rowman and Littlefield Publishers.
- Caplan, A.L. (1992). *If I Were a Rich Man Could I Buy a Pancreas? And Other Essays on the Ethics of Care*, Bloomington and Indianapolis: Indiana University Press.
- Carbonell, V. (2013). *De dicto* desires and morality as fetish. *Philosophical Studies: An*

- International Journal for Philosophy in the Analytic Tradition*, 163(2), 459-177.
- Code, L. (1991). *What Can She Know? Feminist Theory and Construction of Knowledge*, Ithaca: Cornell University Press.
- Connolly, K. (2017). Perceptual Learning. In Zalta, E. N. (Ed.), *The Stanford Encyclopedia of Philosophy*, Summer 2017 Edition. URL = <<https://plato.stanford.edu/archives/win2016/entries/perceptual-learning/>>.
- Collins, P.H. (2004). Learning from the Outsider Within: The Sociological Significance of Black Feminist Thought. In Harding, S.G. (Ed.), *The Feminist Standpoint Theory Reader: Intellectual and Political Controversies*, 103-126. Routledge.
- Cova, F., García, A., & Sennwald, V. (In prep.). *Is What I'm Feeling Really Genuine? Fiction vs. Reality*.
- Cowan, R. (2013a). Clarifying Ethical Intuitionism. *European Journal of Philosophy*.
- Cowan, R. (2013b). Perceptual Intuitionism. *Philosophy and Phenomenological Research*.
- Crosthwaite, J. (1995). Moral Expertise: A Problem in the Professional Ethics of Professional Ethicists. *Bioethics*, 9(4), 361-379.
- Dennett, D. (1969). The Nature of Images and the Introspective Trap. Reprinted in *Content and Consciousness*, 132-146. London: Routledge & Kegan Paul.
- Descartes, (1641). *Meditations on First Philosophy*. In the version by Jonathan Bennett presented at www.earlymoderntexts.com, last amended April 2007, retrieved March 2018.
- Döring, S. (2007). Seeing What to Do: Affective Perception and Rational Motivation.

- Dialectica*, 61(3), 363-394.
- Driver, J. (2006). Autonomy and the asymmetry problem for moral expertise. *Philosophical Studies*, 128(3), 619-644.
- Driver, J. (2013). Moral Expertise: Judgment, Practice, and Analysis. *Social Philosophy and Policy*, 30(1-2), 280-296.
- Elridge, R. (1989): *On Moral Personhood: Philosophy, Literature, Criticism, and Self-Understanding*. Chicago: University of Chicago Press.
- Enoch, D. (2014). A Defense of Moral Deference. *Journal of Philosophy*, 111, 1–30.
- Føllesdal, A. (2004). The philosopher as coach. In Kurz-Milcke, E. & Gigerenzer, G. (Eds.), *Experts in Science and Society*. New York: Kluwer.
- Foot, P. (1967). The Problem of Abortion and the Doctrine of Double Effect. *Oxford Review*, 5.
- Fricker, Miranda (1998). Rational authority and social power: Towards a truly social epistemology. *Proceedings of the Aristotelian Society* 98 (2):159–177.
- Friend, S. (2016). Fiction and Emotion. In Kind, A. (Ed.), *The Routledge Handbook of Philosophy of Imagination*. 217-230.
- Gardner, J. (1978). Moral Fiction. In *Moral Fiction*. New York: Basic Books.
- Gendler, T. (2016). Imagination. In Zalta, E. N. (Ed.), *The Stanford Encyclopedia of Philosophy*, Winter 2016 Edition. URL = <<https://plato.stanford.edu/archives/win2016/entries/imagination>>.
- Goldie, P. (2007). Seeing What is the Kind Thing to Do: Perception and Emotion in

- Morality. *Dialectica*, 2007, 347-361.
- Grasswick, H. (2016). Feminist Social Epistemology. In Zalta E.N. (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition). URL = <<https://plato.stanford.edu/archives/win2016/entries/feminist-social-epistemology/>>
- Gregory, M. (2009). *Shaped by Stories: The Ethical Power of Narratives*. Notre Dame, IN: University of Notre Dame Press.
- Grundmann, T. (2010). Some hope for intuitions: A reply to Weinberg. *Philosophical Psychology*, 23, 481-509.
- Harding, S.G. (Ed.) (2004). *The Feminist Standpoint Theory Reader: Intellectual and Political Controversies*. Routledge.
- Hartsock, N. (1983). "The Feminist Standpoint: Developing the Ground for a Specifically Feminist Historical Materialism," *Discovering Reality: Feminist Perspectives on Epistemology, Metaphysics, Methodology, and the Philosophy of Science*, Sandra Harding and Merrill Hintikka (eds.), Dordrecht: D. Reidel, 283–310.
- Heath, J. (2015, May 30). The Problem of Me Studies. Retrieved January, 2016, from <http://induecourse.ca/the-problem-of-me-studies/>
- Hills, A. (2009). Moral Testimony and Moral Epistemology. *Ethics*, 120(1), 94-127.
- Hoffman, F. (2010). Intuitions, concepts, and imagination. *Philosophical Psychology*, 23, 529-546.

- Hopkins, R. (2007). What is Wrong With Moral Testimony? *Philosophy and Phenomenological Research*, 74(3), 611–634.
- Howell, R. J. (2014). Google morals, virtue, and the asymmetry of deference. *Noûs*, 48(3), 389-415.
- Jaggar, A. (2004). Feminist Politics and Epistemology. In Harding, S.G. (Ed.), *The Feminist Standpoint Theory Reader*, 55-56. Routledge.
- Johnson, M. (2016). Moral Imagination. In Kind, A. (Ed.), *The Routledge Handbook of Philosophy of Imagination*, 355-367.
- Johnson King, Z. (FC). Accidentally Doing the Right Thing, forthcoming in *Philosophy and Phenomenological Research*.
- Jones, K. (1999). Second-Hand Moral Knowledge. *The Journal of Philosophy*, 96(2), 55-78.
- Jones, K., & Schroeter, F. (2012). Moral Expertise. *Analyse & Kritik*, 34(2), 217-230.
- Jones, K. & Schroeter, F. (2018). Moral Expertise. In McPherson, T. & Plunkett, D. (Eds.), *The Routledge Handbook of Metaethics*, 459-471.
- Kant, I. (1785, repre. 1998). *Groundwork of the Metaphysics of Morals*. Trans. Mary Gregor. Cambridge, UK: Cambridge University Press.
- Kind, A. (forthcoming, final MS cited). Imaginative Vividness. *Journal of the American Philosophical Association*.
- Le Guin, U.K. (1973). *The Ones Who Walk Away from Omelas*.
- Lillehammer, H. (1997). Smith on moral fetishism. *Analysis*, 57(3), 187-195.

- Lord, E. (forthcoming, 2017 MS cited). How to Learn About Aesthetics and Morality Through Acquaintance and Deference. *Oxford Studies in Metaethics*.
- Lord, E. (2016). *Syllabus for Phil-002: Ethics*. Fall 2016, University of Pennsylvania (USA).
- Markovits, J. (2010). Acting for the Right Reasons. *The Philosophical Review*, 119(2), 201-242.
- McBrayer, J.P. (2010a). Moral Perception and the Causal Objection. *Ratio*, 23(3), 291-307.
- McBrayer, J.P. (2010b). A Limited Defense of Moral Perception. *Philosophical Studies*, 149(3), 305-320.
- McGrath, S. (2011, MS cited). Moral Knowledge and Experience. *Oxford Studies in Metaethics*, 6.
- McGrath, S. (2008). Moral Disagreement and Moral Expertise. *Oxford Studies in Metaethics*, 3, 87-106.
- McGrath, S. (2009). The Puzzle of Pure Moral Deference. *Philosophical Perspectives* 23: 321–344.
- McGrath, S. (2011a, MS cited). Moral Knowledge and Experience. *Oxford Studies in Metaethics*, 6.
- McGrath, S. (2011b). Skepticism about Moral Expertise as a Puzzle for Moral Realism. *Journal of Philosophy* 108 (3): 111-137.
- McGrath, S. (2011c, MS cited). Normative Ethics, Conversion, and Pictures as Tools

- of Moral Persuasion. In Mark Timmons (Ed.) *Oxford Studies in Normative Ethics*, 268-96. Oxford: Oxford University Press.
- McNaughton, D., & Rawling, P. (2018). Motivating Reasons and Normative Reasons. In (Ed. Daniel Starr), *The Oxford Handbook of Reasons and Normativity*: Oxford University Press.
- Mills, C. (2017). Ideology. In Kidd, I.J., Medina, J., & Pohlhaus, G. (Eds.), *Routledge Handbook of Epistemic Injustice* (100-111). Taylor & Francis.
- Milona, M. (MS). *Armchair Moral Knowledge and the Perceptual Analogy*.
- Milona, M. (2017). Intellect versus Affect: Finding Leverage in an Old Debate. *Philosophical Studies* 174(9), 2251-2276.
- Milona, M. (2016). Taking the Perceptual Analogy Seriously. *Ethical Theory and Moral Practice*, 19(4), 897-915.
- Mill, J.S. (1998). *Utilitarianism*. Oxford: Oxford University Press.
- Moore, A. (2006). Maxims and Thick Ethical Concepts. *Ratio*, XIX, 129-147.
- Moselle, C. (2015). *The Wolfpack*. Dir. Crystal Moselle. Kotva Films; Verisimilitude, 2015. Documentary.
- Naravez, D. & Lapsley, D. (2005). The Psychological Foundations of Everyday Morality and Moral Expertise. In D. Lapsley and C. Power (Eds.), *Character Psychology and Character Education*, 140-165. Notre Dame: IN: University of Notre Dame Press.
- Nussbaum, M. (1990). *Love's Knowledge: Essays on Philosophy and Literature*. Oxford:

Oxford University Press.

Raftopoulos, A. & Zeimbekis, J. (2015). The Cognitive Penetrability of Perception: An Overview. In Raftopoulos, A. & Zeimbekis, J. (Eds.), *The Cognitive Penetrability of Perception: New Philosophical Perspectives*, 1-56. Oxford University Press.

Raz, J. (1975). *Practical Reasoning and Norms*. London: Hutchinson & Co., reprinted, Oxford University Press, 1990 and 1999.

Roeser, S. (2011). *Moral Emotions and Intuitions*. New York: Palgrave Macmillan.

Scanlon, T. M. (1998). *What We Owe to Each Other*. Cambridge, MA: Belknap Press of Harvard University Press.

Schroeder, M. (2014). *Explaining the Reasons We Share: Explanation and Expression in Ethics, Volume 1*. Oxford: Oxford University Press.

Schwitzgebel, E. & Rust, J. (2016). The Moral Behaviour of Ethicists. In Systma, J. & Buckwalter, W. (Eds.), *The Moral Behaviour of Ethicists*. Wiley-Blackwell.

Schwitzgebel, E. & Rust, J. (2014). The Moral Behaviour of Ethicists and the Power of Reason. *Advances in Experimental Moral Psychology*.

Schwitzgebel, E. & Rust, J. (2013). Ethicists' and Non-Ethicists' Responsiveness to Student Emails: Relationships among Expressed Normative Attitude, Self-Described Behaviour, and Experimentally Observed Behaviour. *Metaphilosophy*, 44, 350-371.

Schwitzgebel, E. & Cushman, F. (2015). Philosophers' Biased Judgments Persist Despite Training, Expertise, and Reflection. *Cognition*, 141, 127-137.

- Schwitzgebel, E. & Cushman, F. (2012). Expertise in Moral Reasoning? Order Effects on Moral Judgment in Professional Philosophers and Non-Philosophers. *Mind & Language*, 27, 135-153.
- Sepielli, A. (2016). Moral uncertainty and fetishistic motivation. *Philosophical Studies*, 173(11), 2951-2968.
- Siegel, S. (2006). Which Properties Are Represented in Perception? In Gendler, T. & Hawthorne, J. (Eds.), *Perceptual Experience*, 481-503. Oxford: Oxford University Press.
- Siegel, S. *The Contents of Visual Experience*. Oxford: Oxford University Press.
- Siewert, C.P. (1998). *The Significance of Consciousness*. Princeton, NJ: Princeton University Press.
- Singer, P. (1972a). Moral Experts. *Analysis*, 32(4), 115-117.
- Singer, P. (1972b). Famine Affluence, and Morality. *Philosophy and Public Affairs*, 1(3). 229-243.
- Singer, P. & Wells, D. (1984). *The Reproductive Revolution*, Oxford: Oxford University Press.
- Singh, K. (2018). *Moral Worth, Credit, and Non-Accidentality*. Workshop presentation, Washington University Workshop in Ethics, St. Louis, Missouri, USA, October 13 2018.
- Sliwa, P. (2012). In defense of moral testimony. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 158(2), 175-195.

- Sliwa, P. (2016). Moral Worth and Moral Knowledge. *Philosophy and Phenomenological Research*, 93(2), 393-418.
- Sliwa, P. (2017). Moral Understanding as Knowing Right from Wrong. *Ethics*, 127, 521-552.
- Smith, M. (1994). *The moral problem*. Oxford: Blackwell.
- Smith, M. (2000). Moral Realism. Reprinted in his *Ethics and the A Priori* (Cambridge: Cambridge University Press, 2004): 181-207.
- Sreenivasan, G. (2015). A plea for moral deference. *Etica & Politica*, 17(2), 41-59.
- Thomson, J. (1971). A Defense of Abortion. *Philosophy & Public Affairs*, 1(1), 47-66.
- Tooley, M. (1972). Abortion and Infanticide. *Philosophy & Public Affairs*, 2(1), 37-65.
- Tropman, E. (2018). Intuitionism in Moral Epistemology. In McPherson, T. and Plunkett, D. (Eds.), *The Routledge Handbook of Metaethics*, 472-483.
- Van Leewuwen, N. (2016). Imagination and Action. In K., Amy (Ed.), *The Routledge Handbook of Philosophy of Imagination*, 286-299.
- Weinberg, J., Gonnerman, C., Buckner, C., & Alexander, J. (2010). Are philosophers expert intuiters? *Philosophical Psychology*, 23, 331-355.
- Werner, P. (2016). Moral Perception and the Contents of Experience. *Journal of Moral Philosophy*, 13, 294-317.
- West R.F., Meserve R.J., Stanovich K.E. (2012). Cognitive sophistication does not attenuate the bias blind spot. *Journal of Personality and Social Psychology*, 103(3), 506-519.

- Williamson, T. (2011). On intuitional stability: The clear, the strong, and the paradigmatic. *Cognition*, 115, 491-503.
- Wright, J. (2010). On intuitional stability: The clear, the strong, and the paradigmatic. *Cognition*, 115, 491-503.
- Wylie, A. (2003). Why standpoint matters. In Figueroa, R. & Harding, S.G. (Eds.), *Science and Other Cultures: Issues in Philosophies of Science and Technology*, 26-48. Routledge.
- Yetter Chappell, R. (MS). *The Right Wrong-Makers*. Draft as of August 18, 2018.
- Zagzebski, L. (2003). Emotion and Moral Judgment. *Philosophy and Phenomenological Research*, 66(1), 104-124.