

Against Moral Deference

Adina Cristina Covaci

Submitted in accordance with the requirements for the
degree of Doctor of Philosophy

The University of Leeds

School of Philosophy, Religion and History of Science

September 2019

The candidate confirms that the work submitted is her own and that appropriate credit has been given where reference has been made to the work of others.

This copy has been supplied on the understanding that it is copyright material and that no quotation from the thesis may be published without proper acknowledgement.

The right of Adina Cristina Covaci to be identified as Author of this work has been asserted by Adina Cristina Covaci in accordance with the Copyright, Designs and Patents Act 1988.

Acknowledgments

I want to thank my supervisor, Pekka Väyrynen, for his priceless philosophical guidance. I am grateful for the patience, promptness, and diligence with which he has read and commented on so many drafts of each chapter of this thesis. This thesis would not have been possible without him.

Thank you very much to Alison Hills and Daniel Elstein for giving up their time to be my examiners.

I am also very thankful for the financial support received from the University of Leeds, for the first 3 years of my PhD, and from the Society for Applied Philosophy, for the 4th year.

A big thank you to all the people at Leeds who listened to me talk about this thesis, read many bits of it, and helped to make it better. Thanks also to the Leeds postgraduate community for the support and the fun.

I am incredibly lucky to have met Alison Toop, Bryan Ross, David Heering, Emily Paul, Gary Mullen, Pei-Lung Cheng, and Will Gamester during my time at Leeds. I honestly do not think I could have done this without all of you.

Special thanks to Gary Mullen, who was there with a pep talk whenever I needed it.

Finally, I am most thankful for my family, especially my parents, Ana and Marcel Covaci, whose unwavering support made all of this happen. I dedicate this thesis to them.

Abstract

Is there a problem if you form the belief that capital punishment is morally wrong by deferring to a reliable moral expert? While deferring to your professor about facts concerning physics seems fine, deferring about the morality of capital punishment triggers negative intuitions. In this thesis, I examine these intuitions and investigate whether there are any non-epistemic reasons to not defer about moral matters. I construct and defend a new variety of *moral deference pessimism*, the view that there is something problematic about forming and sustaining moral beliefs, or about acting, on the basis of moral testimony. My account proposes that *recurrent moral deference*, i.e. moral deference that happens repeatedly, is *pro tanto* bad insofar as, and to the extent that, it interferes with the exercise and development of our capacity for practical deliberation. This interference occurs as instances of practical deliberation are being replaced with deference. Thus, when we defer, we do not exercise and do not develop our capacity for practical deliberation. This is *pro tanto* bad because this capacity has instrumental and extrinsic final value. My investigation starts with moral deference, but my practical deliberation view is able to offer a more comprehensive account, which covers other kinds of deference that seem suspicious, such as prudential and aesthetic deference. As such, this project aims to provide a systematic account of the *pro tanto* non-epistemic badness of deference, that is in broad accord with our intuitions, both in morality and beyond.

Table of Contents

Introduction	8
Chapter I: Meet Moral Deference	15
1. What Is Moral Deference?	16
2. The Debate	24
2.1. Pure and Impure Moral Deference.....	26
2.2. The Right Kind of Bad Deference and The Wrong Kind of Bad Deference	29
2.3. There Is Something about Moral Deference.....	31
2.3.1. How Problematic (or Unproblematic) Is Moral Deference? ..	33
2.3.2. When Is Moral Deference Problematic (or Unproblematic)?	36
2.3.3. Why Is Moral Deference Problematic (or Unproblematic?) ..	39
3. Moral Deference Pessimism and Moral Deference Optimism.....	42
4. Concluding Remarks	44
Chapter II: Foes of Moral Deference	45
1. A General Problem.....	46
2. A Way Forward: Recurrent Moral Deference	51
3. The Foes of Moral Deference.....	55
3.1. The Moral Understanding Strategy	55
3.1.1. The Moral Understanding Strategy: The Strong Reading	60
3.1.2. The Moral Understanding Strategy: The Weak Reading	66
3.2. The Virtue-Based Approach.....	70
3.3. The Acquaintance View.....	75
3.4. The Authenticity Argument.....	80
4. Concluding remarks.....	86
Chapter III: Friends of Moral Deference	88
1. The Negative Strategy	89
2. The Positive Strategy.....	94
2.1. The Coping Argument.....	94
2.2. The No Significant Costs Argument	96

2.3.	The Epistemic Injustice Argument.....	98
2.4.	The Close Relationships Argument.....	100
3.	Responding to the Positive Strategy	102
4.	An Interesting Pessimism?	103
5.	Appendix: Optimism outside Moral Deference Optimism.....	105
6.	Concluding Remarks	109
Chapter IV: The Real Problem with Moral Deference		111
1.	Moral Deference and Practical Deliberation	112
1.1.	The Value Thesis	115
1.2.	The Interference Thesis.....	121
2.	Objections.....	128
2.1.	The Overgeneralisation Worry	128
2.2.	Another Overgeneralisation Worry	130
2.3.	Does Practical Deliberation Matter for Moral <i>Beliefs</i> ?.....	132
2.4.	Is Recurrent Deference Truly Incompatible with Deliberation?.	133
2.4.1.	Moral Deference Might Be Recommended by Practical Deliberation.....	133
2.4.2.	Deferring in One Domain, Deliberating in Others	137
3.	Concluding remarks.....	138
Chapter V: The Asymmetry of Deference		142
1.	The Domain-Based Asymmetry Thesis.....	143
2.	Against the Domain-Based Asymmetry Thesis.....	146
3.	A New Basis for the Asymmetry Thesis	157
4.	The Asymmetry Thesis Revisited	160
4.1.	Theoretical Deference.....	160
4.2.	Prudential Deference	163
4.3.	Etiquette Deference	165
5.	Objections: Aesthetic Deference.....	168
6.	Concluding remarks.....	176
Chapter VI: Moral Deference and Metaethics		177
1.	The Oddness of Moral Deference as an Objection to Moral Realism	178

2. Metaethical Explanations for the Oddness of Moral Deference.....	182
3. Metaethical Explanations for the Oddness of Moral Deference: A Response.....	187
4. An Objection.....	195
5. Concluding remarks.....	196
Conclusion	198
Bibliography	200

Introduction

There is a Romanian saying that my mother jokingly brings up whenever I go against her advice and it turns out that she was right: do not expect things to end splendidly, when you do not listen to anybody (it does not translate that well). In colloquial terms, this thesis is about whether my mother is right, and we should be listening to what other people tell us.¹ In philosophical terms, this thesis investigates the non-epistemic status of forming beliefs and acting on the basis of another person's judgement. In other words, here I seek to find out whether there are any non-epistemic reasons to not defer to others' verbal or non-verbal testimony.

It is already a cliché to say that it is a cliché that we all rely on testimony: a lot of what we know about ourselves and about the world comes from other people. The practice is widespread, and it has been the object of study of many epistemologists for quite some time now. Because of that work (e.g. Fricker and Cooper, 1987; Lackey and Sosa, 2006; Pritchard, 2004), I will take as an assumption that testimony, in general, is an epistemically legitimate way of obtaining knowledge. So I will be neutral on specific issues regarding the epistemology of testimony, such as whether testimony offers us justification by itself or we need to appeal to other sources of knowledge (perception, memory, inference) — my arguments are indifferent to whether you are an anti-reductionist (the former view) or a reductionist (the latter view) about testimony. I will also assume that the relevant conditions that allow for an epistemically appropriate production or transmission of knowledge can, at least in principle, be met (that there are experts to whom we can defer, and we can identify them with confidence; that they can be reliable and trustworthy; that we can be epistemically safe and responsible in taking testimony). So I will bracket any epistemic questions, because I want to focus on the *non*-epistemic status of deference. However, if someone is not convinced that the epistemology of testimony can work, my investigation can be considered under a conditional clause: *if* the epistemology of testimony is right, are there any non-epistemic reasons to not defer to others?

¹ Of course, we could 'listen to what others tell us' by actually deferring to them or just by letting our own reasoning be guided by what they say, as we will see below.

Take the following case:

Suppose those wizards at Google come out with a new app: Google Morals. No longer will we find ourselves lost in the moral metropolis. When faced with a moral quandry [*sic*] or deep ethical question we can type a query and the answer comes forthwith. Next time I am weighing the value of a tasty steak against the disvalue of animal suffering, I'll know what to do. Never again will I be paralyzed by the prospect of pushing that fat man onto the trolley tracks to prevent five innocents from being killed. I'll just Google it. Again I find myself dependent on Google for my beliefs, but in this case it seems to many, myself included, that this is not a good way to go. *There seems to be something wrong with using Google Morals. But what is it?* (Howell, 2014, p. 390, my italics)

My research also starts from this intuition, and aims to provide answers to the following questions: what does it mean that there is 'something wrong' with deference in certain cases? What explains our intuitions? Is there genuinely something wrong with (at least some cases of) deference? If there is, what kind of problem is it? If we can find a problem, does it apply only to certain types of deference? That is to say, is moral deference uniquely problematic, or particularly so, by comparison to other types of deference? What can we say about cases of non-moral deference? In short, my objective is to offer a systematic account of what, if anything, makes (at least some cases of) deference non-epistemically problematic.

I will explore these issues by focusing particularly on *moral deference*, and then I will follow the implications this investigation has for deference outside of the moral realm. Starting the project by focusing on moral deference makes sense for three main reasons. First, it is the paradigmatic type of deference which raises questions about whether it is problematic to adopt other people's testimony. Secondly, it seems to trigger particularly strong intuitions from people, which is a useful starting point. And, finally, I choose it as my main subject also because, unoriginally, I consider morality to be very important in our lives as human beings. And it is quite intuitively plausible that some people are better than others at figuring out how to best act morally or what moral

beliefs we should hold.² So the question which motivated me was the following: given how important morality is for us, should we just take the testimony of such people about moral matters — perhaps all the time, if possible? After all, they know (or are) better! Or are there reasons not to do that?

My answer is that although there may be no reason to refuse testimony, no matter what, all the time, there are indeed reasons not to defer about moral matters *recurrently*. I will put forward a view which aims to show that recurrent moral deference is *pro tanto* bad. But my approach to the topic will end up revealing results that apply beyond moral deference. I will argue that moral deference is not uniquely special in any substantive way and then use this as a starting point to explore the non-epistemic status of deference more generally. I will argue that whatever non-epistemic reasons we might have to not defer in the moral case seem to apply to many other cases of deference as well. I will illustrate how my argument against moral deference extends to certain other types of deference and how it can provide a principled way of distinguishing between problematic and unproblematic cases of moral and other types of deference.

Here is how I plan to do this. In chapter I, I introduce the debate on moral deference. I begin by showing how the debate starts with looking at various instances of deference, moral and non-moral, that seem to trigger very different intuitions. Subsequently, I do an analysis of the concept of deference, by examining the different possible ways one could spell it out. This is aimed at precisifying the target of the discussion and at zooming in on the exact notion I will examine in later chapters. I also describe the two main sides of the debate, namely *moral deference pessimism* and *moral deference optimism*, and offer my best interpretations of their central claims. The aim of this chapter is to clarify some of the ambiguities that appear in the existing literature and to pave the way for investigating whether there truly is something problematic about moral deference.

In chapter II, I present and evaluate moral deference pessimism. I argue that existing pessimist views about moral deference, although deeply insightful, have failed to give a fully adequate explanation of the problematic character of

² This is not entirely uncontroversial. Although I take this claim as an assumption, I offer a discussion in chapter I on the epistemic arguments against moral deference. The idea also resurfaces in the metaethical context — see chapter VI.

deference. This will further motivate the need for a novel position that can properly achieve that task, and thus make some space in the literature for my own version of pessimism which I develop in chapter IV. I will discuss and reject the views that I take to be the most important and the most developed in the literature, namely what I call *the moral understanding strategy*, *the virtue-based approach*, *the acquaintance view*, and *the authenticity argument*. In this chapter I also argue that we need to shift our focus when it comes to what kind of instances of moral deference we target. Investigating isolated or individual cases of moral deference (e.g. I defer to you one time about whether eating meat is morally permissible) is not fruitful. So I will propose to investigate *recurrent moral deference*, i.e. moral deference which happens repetitively. This will give us a better chance at assessing moral deference in itself or *per se*, which also lays the foundations for my own version of moral deference pessimism.

In chapter III, I present and evaluate moral deference optimism. I argue that the existing optimist accounts are not successful in showing that pessimism is wrong. Since my own version of pessimism was not on the market when these views were developed, I will discuss them in the existing landscape first — i.e. existing views of optimism versus existing views of pessimism. I will show that, as they stand, they are in fact compatible and do not seem to be in conflict. Subsequently, I introduce my own version of pessimism, the beginnings of which are developed in chapter II, and show that optimism needs to be developed more in order to actually challenge it.

In chapter IV, I formulate a novel explanation of what makes recurrent moral deference problematic. I argue that recurrent moral deference is *pro tanto* bad insofar as and because it interferes with the exercise and development of our capacity for practical deliberation. My argument consists of two main claims. The first one, *the value thesis*, trades upon the significance of the capacity for practical deliberation in our lives, i.e. its instrumental and (extrinsic) final value. The second claim, *the interference thesis*, shows that recurrent moral deference interferes with the exercise and the development of the capacity for practical deliberation. I hold that the two theses, together with the idea that it is *pro tanto* bad to interfere with something of value, point to the conclusion that recurrent moral deference is *pro tanto* bad. Of all the existing accounts, my practical deliberation argument, I will hold, is best suited to explain the widest range of

cases, correcting for the over- and undergeneralisation that have plagued previous views.

In chapter V, I discuss *the asymmetry thesis*, the claim that there is a contrast between various kinds of deference, such that some are problematic, but others are not. This claim is motivated by the different intuitions that are triggered when we consider deferring to our physics teacher or accountant versus deferring to an ethics teacher or a moral expert. Intuitively, the former kind of deference is not problematic while the latter is. In this chapter I first examine how the asymmetry thesis has been discussed in the literature. Seemingly, philosophers have mainly suggested that we can distinguish between problematic and unproblematic deference by looking at which domains of discourse the matters we defer about belong. More specifically, the existing ways of accounting for the asymmetry propose mapping the problematic-unproblematic distinction onto the moral-non-moral or, alternatively, normative-non-normative distinction. However, by examining different examples of deference, I argue that this is not the right way to look at the asymmetry because it does not capture our intuitions correctly. Instead, I propose that my practical deliberation account can offer a principled way to distinguish between which cases of deference are problematic and which are not, in a way that is in broad accordance with our intuitions. Given that moral deference seems to pattern with other kinds of deference, the appeal to practical deliberation can systematically explain why various instances of deference, be them moral or non-moral, are problematic. This makes my position the most comprehensive one in the literature, as it can do something that the rival views cannot do or have a harder time doing. Namely, my argument provides a more unified explanation for the seemingly problematic character of other types of deference, e.g. the aesthetic and the prudential, and it can also give the right predictions in the various normative and non-normative cases, e.g. legal, etiquette.

Finally, in chapter VI, I investigate moral deference under different metaethical views. Although first-order, axiological, explanations for the problematic character of moral deference are available, there is a question regarding the possibility and plausibility of a metaethical explanation. To answer that, I explore whether different metaethical views can attempt to explain our attitudes towards moral deference and whether our intuitions about moral

deference put any pressure on the metaethical commitments that we might have. At the same time, there is the question of whether the first-order explanation that we have for the problematic character of moral deference is incompatible with any metaethical view. I argue that a second-order metaethical explanation for the negative intuitions triggered by moral deference cannot capture what is bad about moral deference because there is a first-order, axiological problem with it. Moral deference is *pro tanto* bad vis-à-vis some value that it diminishes or makes us lose, and metaethics cannot tell us why that is the case, since it operates at a different level. I also show that my practical deliberation argument holds robustly across a range of metaethical views.

To end this introduction, I want to say why this thesis is worthwhile. I take this project to be significant and interesting precisely because taking testimony from other people is a practice that is so central to our lives. We are social creatures who live with each other and share many things, including epistemic resources. So, broadly speaking, not only death and taxes are certain or inevitable, but taking testimony from others too. Moreover, many of us care a lot about morality. We care about why and how people come to believe certain things and act in certain ways, as this relates to aspects of autonomy, practical deliberation, motivation, virtue, moral worth and blameworthiness, among other things. Some might even think that all of this makes morality special in a sense, such that moral deference itself is special. It is thus important to see what place there is for moral deference within our moral agency, and whether it is indeed special in any way. And whatever the answer, we need to see what can be said about deference beyond the moral case. However, I believe that the previous attempts to analyse the normative status of deference have not quite managed to get to the heart of the problem. On the one hand, although they draw attention to important features of the phenomenon, the accounts which argue in favour of moral deference do not manage to explain away all the problematic cases. On the other hand, those who argue against moral deference fail to properly cover all the problematic cases and to explain what is amiss about moral deference. While they provide some illuminating commentary on moral deference, they ultimately do not offer the full story. So I believe a better explanation is needed and I hope to make it convincing that I present a worthwhile one.

I embark on this project recognising that the philosophical way of discussing deference seems, at least at times, artificial and unrealistic. The almost clinical definitions and distinctions that I will put forward might seem alien when compared to our everyday moral (or aesthetic, prudential etc.) reasoning and decision-making. Perhaps you might say that no one defers in the clinical way philosophers portray it. Or you might even think that no one defers at all — we all just use other people's opinion as food for thought and then do our reasoning ourselves. However, I do not believe this to be true. While some examples used in the literature are indeed unrealistic, we will see later that others easily strike us as familiar. One of my aims is to bring the discussion closer to reality and show that the clinical way is not the only way to talk about deference.

As for the thought that maybe deference does not happen at all, it is, of course, difficult to say with certainty; after all, we can only rely on people's testimonies as to whether or how they take testimony. But by reflecting on our own experience, we are likely to find some examples of actual deference. Think about whether you have ever caught yourself at least *prima facie* adopting a view or a verdict in a controversial moral case just because it comes from someone from the same political or moral side as you, whom you see as a relative expert. This does not seem so implausible, and I take it that this is deference in its natural habitat. You might go on and try your best to deliberate about the issue yourself, in which case it ceases to be deference. But you might not. This is not an unrealistic example at all, I contend. So, I will offer some clinical examinations in what follows because our everyday moral reasoning and the concepts we use can get messy, and so do not offer, as such, a good starting point. But what I will say very much applies to our everyday lives, which I think makes this project valuable.

Chapter I: Meet Moral Deference

The chemical formula of sulphuric acid is H_2SO_4 . The world's smallest mammal is Kitti's hog-nosed bat. The speed of light is 300,000 kilometres per second. The overwhelming majority of people who know these facts, and many more of this kind, know them because other people have told them so; in other words, by taking testimony or by deferring to others. A lot of information about geography, chemistry, medicine, physics, or history has been transmitted through testimony, and we have no problem in considering we have knowledge of these matters. We take testimony to be a reliable way of acquiring knowledge and "our dependence on testimony is as deep as it is ubiquitous" (Lackey, 2006, p. 1). Our lives would be very much impoverished, and probably a lot more difficult to navigate, if we excluded testimony as a legitimate source of knowledge, and if we thought less of the people who rely on it (i.e. all of us!). Of course, as all sources of knowledge, testimony is not infallible. What I have just said does not imply that we should blindly believe anything that people tell us. It is not epistemically good to gullibly form beliefs on the basis testimony, if we know that the source is unreliable or not worthy of trust.³ However, if no obstacles of this sort exist, as a rule of thumb, it seems unproblematic, both epistemically and (at least *prima facie*) non-epistemically, to defer to other people's testimony.

Things are different, however, when we consider moral matters. Take the following examples: I know that abortion is morally permissible because my ethics professor told me so. John believes that the war is just because his wife believes that. My neighbour does not eat meat because her friend told her it is immoral. The way most of us feel about such cases and the people involved is very different from how we feel about someone who has knowledge of the properties of sulphuric acid from her chemistry teacher.⁴ The intuitions shift: there seems to be something "illegitimate" (Hopkins, 2007, p. 617),

³ See Adler (2017) for an overview and Fricker (1994) against gullibility.

⁴ The issue of moral deference is discussed here, and everywhere else in the debate, as relating to adult moral agents. Everyone accepts that it is in no way bad for children to defer, as this is how they learn about morality and become moral agents themselves (Hills, 2009, p. 94; Howell, 2014, p. 390). Moral deference can, and probably should, unproblematically be part of children's moral education. I will take this as an assumption as well.

“problematic” (McGrath, 2009, p. 322), “sub-optimal” (Howell, 2014, p. 390), or “fishy” (Enoch, 2014, p. 229) about moral beliefs and actions that are based on moral testimony.⁵

This is how the debate on this topic starts: with a perceived difference in intuitions about various instances of deference.⁶ The task is then to explain these intuitions. Are they warranted? Do they show anything about moral deference? What do they show? Or, alternatively, can we explain them away? These issues can be categorised under the umbrella of *the problem of moral deference*. In this chapter I introduce this problem. I discuss the concept of moral deference and why it has been the subject of philosophical debate. The aim is to clarify this notion and give a broad overview of the kind of issues to which it gives rise, as well as an outline of the landscape of the debate on moral deference.

1. What Is Moral Deference?

Let me start with an example of deference:

MEDICAL SHOW. Maria and Jennifer are watching a medical TV show. After the show is over, Maria comments on the story line and says that it is not good because it suggests that a doctor can override autonomous people’s informed medical decisions if they think they are making the wrong choice. Instead, Maria claims, doctors morally ought to respect the informed medical decisions of autonomous people, even if they think they are not acting in their own best interests or they disagree with the choice. Jennifer did not see it like that. Even though she did not have a firm opinion, she thought it was fine for the doctor to override the patient’s decisions because, after all, she did what was best for the patient and saved their life. But she knows Maria is more informed about the ethical side of things, and that her parents are both doctors. She also trusts Maria is a good and reliable person. Because of this, she gives up her own belief and adopts Maria’s moral verdict.⁷

⁵ This shift of intuitions also occurs in other kinds of deference, like aesthetic or prudential deference. But I start with *moral* deference as my study case and I will go on to explore other kinds of deference in chapter V.

⁶ I reserve a discussion regarding why our intuitions seem to operate at a first-order, axiological level for chapter VI. This is because it will be particularly salient to the task I undertake there, namely considering whether there is a metaethical explanation for the oddness of moral deference.

⁷ The trained eye will observe that this does not look like the most common examples in the literature because Maria also gives Jennifer an explanation for the moral judgement that the latter ends up deferring on. Typically, philosophers in this debate use examples where no explanation

Aligning with the existing debate, I define the term ‘deference’ as *the act of adopting a judgement as one’s own, by forming and sustaining a belief or performing an action, solely on the basis of the expertise of the person who provides it (and not due to the reasons that actually ground it or make it correct)*. I will use the term ‘moral deference’ as a local version of this, to refer to the *moral* judgements that are being adopted by a person on the basis of someone’s expertise. It is worth spending some time unpacking this definition.

The simplest form of deference typically involves one person taking the testimony of another person, in the sense of adopting it as one’s own view. A very concise way of defining the philosophical (as opposed to the legal) concept of ‘testimony’ is the following:

when someone tells us p , where p is some statement, and we accept it, then we are forming a testimonially-based belief that p . Testimony in this sense need not be formal testimony in a courtroom; it happens whenever one person tells something to someone else. (Green, 2008)

Testimony can be oral as well as written. But although taking testimony is the most common way of deferring, the concept of deference can cover special cases too. For example, it can cover instances where people may form a moral judgement based on reading someone’s body language and non-verbal reactions rather than speech; based on overhearing someone utter a moral judgement; or with the help of a fMRI from the future which would help one see the moral judgement in another person’s brain (Howell, 2014, p. 394).⁸ These are not examples of testimony, but they are examples of deference. From now on, I will be using testimony as my primary case of deference, but everything I will say can be applied to whatever form of deference one has in mind.

is given to the deferrer about the judgement on which they defer. I explain below the complications that can arise in cases such as the one I give here, but I stand by my example because I think it is a bit less clinical and a little more realistic and plausible (which will become relevant later).

⁸ Not everyone will allow these special cases to count as relevant. Wiland (2017, p. 52), for example, does not, and limits the term “to tellings intended to be accepted as testimony by the speaker in question”. I see no reason to restrict the term, but everything I say applies to however you want to conceive of the extension of deference.

I say that deference is the act of adopting a judgement as one's own, by forming and sustaining a belief or performing an action, *solely* on the basis of the expertise of the person who provides it (and not due to the reasons that actually ground it or make it correct). In the literature, this is called *direct moral deference* (Fletcher, 2016, p. 48). This is meant to isolate deference because, in fact, there are many things one can do with a piece of testimony, and not all of those things will count as deference.⁹ For example, this definition is firstly meant to immediately distinguish deference from acts of rational persuasion. Say that you simply tell me that abortion is morally permissible, without any explanation, and I adopt this judgement as my own; then I defer to you. By contrast, say you tell me that abortion is morally permissible and also why that is so. You tell me that women have the right to control their bodies and that foetuses are, at least in early stages, not persons etc. I *understand* how the reasons support the claim and come to believe that abortion is morally permissible because of that. This is not moral deference; it is just rational persuasion. My new belief has been formed on the basis of the reasons that support it and not solely on your expertise. Your testimony would not figure in an explanation of my belief that abortion is morally permissible (except circumstantially). You *convinced* me of the truth of this moral judgement and there is nothing problematic about forming beliefs in this way or acting on this basis. The definition above is also meant to exclude cases where one uses someone's testimony as food for further thought, as a reason among others, or as corroborating evidence in coming to form and sustain a new belief.

Now some philosophers have allowed forming a belief *partly* on someone's testimony to be called 'deference' (*indirect moral deference*, Fletcher, 2016, p. 48). I think this complicates things unnecessarily. If we bring in other sources, such as inference or perception, then we are muddying the waters. It is not non-epistemically problematic to form beliefs based on those sources of knowledge, so there is no point in looking at them. We want to isolate and zoom in on testimony, and that is why we need to look at cases where only testimony is involved. We can do many things with testimony and there is no need to call them all 'deference'. If we want to ensure "that the relevant characters do not

⁹ See, for example, Nickel's (2001, p. 255) categorisation of the ways in which we can depend of testimony.

hold their beliefs for other reasons” (Lord, 2018, p. 74), so as to not distract our evaluation of moral deference, I think it is just easier to restrict ‘deference’ in the way I do.

I thus want to reserve the term ‘deference’ for when a person adopts and sustains a moral judgement or acts solely on the basis of what another person tells them. For example, I appropriate this judgement because you uttered it, and I consider you to be an expert. Roughly, I believe you are good at getting the right answer and I trust you. So, following Wiland (2017), I will hold that the marker for what counts as deference will be that the expert’s testimony is the *sole* explanation for the newly acquired belief or action performed. Or, as Wiland (2017, p. 53) puts it, “accepting testimony occurs only when the testimonial reasons alone either *do* or *would* make a difference between believing or not believing”.¹⁰ Although in practice, it might be difficult to pin down exactly how much someone’s testimony contributes to someone’s formation of a new belief, in principle we can make the distinction, which I take to be sufficient for positing what instances of taking testimony will count as deference.

This is also the kind of case on which that the literature focuses. Take, for example, one of the most used examples:

Eleanor has always enjoyed eating meat but has recently realized that it raises some moral issues. Rather than thinking further about these, however, she talks to a friend, who tells her that eating meat is wrong. Eleanor knows that her friend is normally trustworthy and reliable, so she believes her and accepts that eating meat is wrong. (Hills, 2009, p. 94)

This is a clear case of moral deference because Eleanor unambiguously adopts her new belief solely because of her friend’s testimony. The fact that her friend

¹⁰ This case can be tricky because your testimony can contribute more or less to my final belief. Wiland (2017, p. 53) considers four scenarios regarding such a situation. One, either your testimony or my non-testimonial reasons are enough for me to form that belief; the belief is overdetermined here. Two, the testimony-based reasons are enough for me to form the belief, but the non-testimonial reasons increase my confidence in it. Three, we have the reverse, where my non-testimonial reasons are enough, but the testimonial reasons give me more confidence. Four, neither your testimony nor my non-testimonial reasons are individually sufficient for me to adopt the new belief, but together they are. To make things easier I will follow Wiland in saying that the testimony needs to make or would make a difference to whether you form the belief (this excludes scenario three, as he remarks).

is a moral expert and has told her that eating meat is wrong is the *only* explanation for why she believes what she believes now.¹¹ Moreover, in this case, no explanation accompanies the claim that eating is wrong.¹² But this need not be the rule (and indeed it probably is not, in real life).

Rather, it is possible for an expert to also tell us the reasons that ground the claim we defer about, like in the MEDICAL SHOW example above. We can thus come to know them. But whether they explain why we have this new belief is a different matter. And I take the difference to be one in *seeing* how the reasons ground the claim. If I see or understand how the reasons ground the claim, then *that* will probably explain why I hold it rather than the testimony itself. If I do not see the relation that holds between claims and reasons (as it is the case when I merely *know* the reasons),¹³ then the testimony will explain why I hold the belief.¹⁴ This is related to why I say in my definition that deference includes belief formation *and* sustaining. Deference sustains a belief if the explanation for the belief continues to be that the expert told me so.¹⁵ If I end up understanding the reasons for the judgement and see how they support it, deference goes out of the picture; it transforms into rational persuasion. The genuine explanation for my belief now ceases to be the expert testimony even though I have initially formed the belief based on it.

It is worth mentioning here, however, that although cases like Eleanor are widespread in the debate¹⁶ and are useful, in a way, to isolate the phenomenon we are interested in (no risk of this being a case of rational

¹¹ Of course, some background conditions need to be in place for this to happen: that she understands the language, that she knows what the concepts involved mean etc. I do not take these as explanations for why she has the new belief.

¹² Lord (2018, p. 74) calls this “deference about a pure assertion”. He chooses to focus on this, as does Fletcher (2016, p. 48). Deference about an impure assertion includes, besides the moral judgement, an explanation for that moral judgement.

¹³ This is based on how Hills (2009) explains understanding, although I do not subscribe to her intellectualised account of understanding. I borrow only the elements described here.

¹⁴ I choose understanding rather than knowledge to mark the line between deference and non-deference inspired by Hills (2009, p. 111, fn. 29; p. 116, fn. 38). The idea is that the belief formed on deference counterfactually depends on the expert’s testimony in a way that a non-deferential belief does not. Compare: the expert tells you that *p* and also the reasons why *p*, but in one case you understand how they connect and in the other one you do not. In the first case, if the explanation for why *p* turned out to be false, then you would stop believe that *p*. By contrast, in the second case you would still go on believing it.

¹⁵ Fritz (2018, p. 121) also mentions that the relevant kind of deference needs to also sustain a belief.

¹⁶ See similar examples used by Hopkins (2007), Howell (2014), Lord (2018), McGrath (2009), Mogensen (2015), and Nickel (2001).

persuasion!), they are intuitively quite bizarre. In the literature, such cases also focus on people who have a moral question and just defer about what the answer is, usually without thinking about it themselves (sometimes whether or not they do think about it is omitted from the case description), like Eleanor above who defers about whether eating meat is wrong. No explanation about why eating meat is wrong is asked for and no explanation is given.

These examples, however, do not seem very plausible or representative of what people actually do. This has made pessimism a target for what I take to be fair criticism from the optimists, e.g. Sliwa (2012, p. 178), Groll and Decker (2014). I take the use of such cases as inadequate, for a number of reasons which I present in chapter II, where I also present my own view on the kind of examples I think we should be using, namely cases of *recurrent moral deference*. Here it suffices to say that using unrealistic and implausible examples hinders our investigation into moral deference, because they easily raise eyebrows. Is it not very peculiar that Eleanor is able to figure out that eating meat might be morally suspicious, yet simply not think about the issue further, and just assume her friend's belief without giving it a second thought? It strikes me that it is. The characters illustrated in the preferred examples of the literature — which are used as intuition pumps and then to ground arguments — present strange moral psychologies.

We are right to be wary of these characters, but not for reasons that have to do with their deference — but rather precisely because of their strange moral psychologies. An insufficient diet of examples, as Groll and Decker (2014, p. 60) call it, can lead us to make incorrect claims and give wrong diagnoses because we are not taking into account the variety that exists within this phenomenon. We will need to consider people who are *involved* with their problem and who perhaps resort to deference only as a last solution. This type of cases can give us insight into what, if anything, is wrong with moral deference *per se*. So, although we want to make sure that we use cases of moral deference where it is clear that one forms a belief or acts *solely* on the basis of testimony, there is no need to make them implausible in the way described above. There is no need to restrict the range of cases we look at to those where there is no explanation (asked for and/or given) accompanying the judgement that is adopted by deference. In this sense, the cases used in the debate thus need improvement.

One last point also worth making in this context is about *advice*, and whether taking advice counts as an instance of deference. I do not take advice to be an instance of deference. I am on the side that sees advice as using someone's testimony only to guide one's reflection and making up one's mind alone in the end (Driver, 2015; Hills, 2009; McGrath, 2009). Taking advice is another thing one can do with a piece of testimony. Thus, taking advice is not problematic on this view and so I will not be concerned with it here.¹⁷

I also define moral deference in relation to a *moral expert*, the person whose judgement the deferrer adopts as their own. In the literature, a moral expert has been described as someone who “very reliably, though not necessarily infallibly, provides correct moral advice in response to moral situations and quandaries” (Cholbi, 2007, p. 325). Moral experts are those who are both “deserving of trust with respect to their moral judgements” and “have greater claim to moral knowledge” (Driver, 2006, p. 625). I follow the view that moral expertise can be diverse (Driver, 2013; Jones and Schroeter, 2012; Miller, 1975). That is, moral expertise can be either practical or theoretical: moral experts can be people who have knowledge of moral principles and are experts at applying them and doing moral reasoning using them;¹⁸ or they can be people who have practical knowledge or experience, such as virtuous people, who have practical wisdom and sensitivity to moral considerations. And, of course, various mixtures are possible. Moral expertise can also be local, relative to a particular area of morality, or global, encompassing all of morality. So the term ‘moral expert’ in my definition includes all these possible kinds of expertise. I am also happy with a weaker sense of the term: we can say someone is a moral expert if they have more knowledge or are better placed to figure out the answer than the deferrer. In other words, someone can be a moral expert *relative* to the deferrer. What specific conditions a person needs to fulfil to be considered a moral expert do

¹⁷ Not everyone agrees with this distinction. Other philosophers hold that taking advice is only a subclass of the deference, i.e. specifically testimony on practical questions (Sliwa, 2012; Wiland, 2000). According to this view, both advice and testimony can be equally problematic or unproblematic. I think this whole discussion is distracting because it seems to turn on how terms such as ‘advice’ or ‘deference’ are defined. The philosophers on the two sides just seem to use ‘advice’ differently. Here it suffices to say that, for me, the relevant cases are those where the sole explanation for the newly acquired belief or action is the testimony of the moral expert. This is what I call deference — but if one conceives of advice as deference about practical questions, then one can think that the notion I use includes advice (in that sense) too.

¹⁸ Singer (1972) has this view, and thinks moral philosophers represent a good example of moral experts. God help us all then.

not concern me here, as my focus is not on the epistemic issues surrounding moral deference. My investigation gets off the ground as long as someone can at least in principle fulfil those conditions and be a moral expert.¹⁹ If moral deference can be epistemically acceptable, my question is about its possible non-epistemic issues. So I will take it that it is possible that moral experts exist.

Moreover, it is not only moral expertise that is diverse. Moral deference — and deference in general — is diverse as well.²⁰ The examples used in the literature are typically very simplistic: I defer to you about whether eating meat is wrong; you defer to me about whether you should join the protests. But deference is not just about adopting full judgements like ‘eating meat is morally impermissible’ or ‘joining the protests is morally required’. For example, say I am trying to decide whether to take the work promotion that will entail me and my family moving to another country. I think about it and I do not know what to do, so I ask my reliable and virtuous friend. But perhaps I do not just need to be told what to do. I might only need help seeing which considerations are morally relevant and which I should take into account in making my decision, e.g. the career of my partner, whether it would upset my children, the importance of professional and financial benefits, the happiness of my family and my own. Groll and Decker (2014, p. 71) call this *relevant reasons variety of testimony*.²¹ Or I might need help weighing these aspects up against each other: which is more important and what should take precedence. Howell (2014, p. 392) calls this *derivative moral deference*, while Groll and Decker (2014, p. 71) call it *interface moral deference*. Wiland (2017, p. 55), for example, takes these instances of deference to be the most realistic and common ones. Of course, I could also be deferring about more fundamental moral values, adopting some values that I have not previously held. Howell (2014, p. 392) calls this *foundational moral deference*. For instances, when I consider the question of vegetarianism, perhaps I start by not

¹⁹ Besides there being moral experts, some further epistemic conditions need to be met for deference to be epistemically safe and responsible. Typically, we should be able to trust the people we ask advice from, and should have good reasons to do so or good reasons not to doubt them. We should not be gullible and just believe anyone or someone whom we have reasons to doubt.

²⁰ Howell (2014, pp. 391-392) has previously done such a survey, and so did Groll and Decker (2014, p. 71) to a shorter extent.

²¹ Fletcher (2016, p. 51) and McGrath (2009, p. 322) seem to think this kind of deference is not problematic if the deferrer and the testifier share a moral framework or sensibility.

believing that animal pain is morally important, but I could defer to someone and adopt that value.²²

A further distinction can also be made regarding the upshot of deference: the result of deference can be the formation of a belief or the performance of an action, in response to testimony. Say Eleanor finds out from her friends that eating meat is immoral so she forms the belief that eating meat is immoral — this is deference in belief or *doxastic deference* (Howell, 2014, p. 391); she may or may not put this belief into practice and become a vegetarian. But it can also be the case that we can take a course of action that follows the testifier's judgement, simply because we take the testifier to be an expert, without assuming that claim as our own belief and introducing it into our own belief system. Howell (2014, p. 391) calls it *active deference*. For example, Tamara, who is a soldier fighting in a war, follows her general's orders to take as prisoners the people she found in a house during a raid because she trusts the general's expertise. Even though she executes the order, she need not form the belief that taking those people as prisoners is morally permissible or obligatory. She need not form any belief at all. Yet she still defers, as she acts solely on the basis of testimony.

I hope that by now the notion of moral deference that is at the centre of this thesis is clearer. This is the basis of the discussion to follow, but some distinctions will be added later in this chapter and then in chapter II. Now I want to move on towards offering a picture of the current debate on moral deference.

2. The Debate

The literature that has focused on this problem can be roughly divided into two sides, according on their stance on moral deference. Taking the labels from Hopkins (2007, p. 613), we have *moral deference pessimists* and *moral deference optimists*. To just give you a hint of where these positions stand, let me say for now that, very roughly, the former have a more negative attitude towards moral deference and the latter a more positive attitude.

However, it is surprisingly difficult to provide a neat characterisation of pessimism and optimism because both come in many shapes and forms, and it

²² Some philosophers take this kind of deference to be the most problematic one (e.g. Howell, 2014, p. 392).

is not very clear how strong or specific their claims are. A position that counts as pessimist under one specification of the view may also turn out to be optimist under a rival specification.²³ Take, for example, Hills' (2009, p. 97) description of the pessimism she defends. The main claim is that:

there are circumstances in which you have no reason to trust moral testimony (in fact, you have reason not to put your trust in it), even if your interlocutor is reliable and trustworthy regarding the matter in question, and you know her to be so.

Hopkins (2007, p. 617) claims that “Pessimism is defined by the claim that it is illegitimate to rely on the word of others in moral matters (unless such reliance is unavoidable)”. On the optimist side, Sliwa (2012, p. 177) argues that “that there is no general problem about moral testimony; in fact, moral testimony is no more problematic than nonmoral testimony.” These claims, although easy to digest and perhaps intuitive, are ambiguous — with respect to what ‘general problem’, ‘problematic’ or ‘illegitimate’ means; what the range of circumstances in which we should not defer is. This matters especially given that they need to be tested against various cases of moral deference, regarding which they might not be able to give a clear verdict.

So I think that it is not sufficient then to say that pessimists think there is a general problem with moral deference and optimists do not, because there are too many moving parts. For the purposes of introducing the debate, I will do some reconstruction work. Thus, I will try to offer here *the best interpretations* of moral deference pessimism and moral deference optimism. The characterisations will be fairly minimal, to capture all possible versions. I will reconstruct the debate by looking at the variables that can change from view to view. Then I will work from there to piece together the best versions of pessimism and optimism that can be extrapolated from the existing literature, with the aim of showing what the core of the conflict between them is. First, I will examine two issues related to moral deference as the target of these views,

²³ One such example is Lord's (2018) view. He takes the disagreement between the two sides to be about whether or not we have an *obligation* not to defer. Since he does not think there is such an obligation, he calls himself an optimist even though he thinks there is something amiss about moral deference. Enoch (2014) also seems to lie at the intersection of optimism and pessimism. He purports to defend moral deference by arguing that it is sometimes morally required. But he also holds that there is something regrettable about it.

to further precisify the phenomenon on which they focus. Subsequently, I will investigate the strength, scope, and variety of the negative claims about moral deference posited in the debate, as well as their negations, under the umbrella of the positive claims about moral deference.

2.1. Pure and Impure Moral Deference²⁴

Take the question regarding the moral permissibility of eating meat. Perhaps I am just starting to reflect on this, but I know someone who is better placed than me at figuring the answer out, and I end up wanting to defer to them. Say I am not sure about some of the relevant empirical aspects of the problem: I do not know how animal farming works, for example what animals are fed and how they are treated. My friend, however, knows all of this and I know she has similar moral sensibilities to mine. So, I defer to her about the moral permissibility of eating meat. Alternatively, say that my friend and I are equally informed about the empirical aspects of the problem, but I still want to defer to her on whether eating meat is morally wrong and whether I should become a vegetarian. According to McGrath (2009) we can call the first instance *impure moral deference*²⁵ and the second *pure moral deference*.²⁶ As the labels suggest, the difference consists in the nature of the content of deference: one concerns exclusively moral information and the other is fundamentally about (morally relevant) non-moral information.

The current literature has focused on *pure* moral deference because, if there is one thing that everyone agrees on, it is that deference about non-moral issues (even if morally relevant) is not problematic. Remember the perceived asymmetry in our intuitions and the examples from the beginning. It seems fine

²⁴ Similar distinctions to the ones I discuss here can be drawn and applied to deference on matters outside of morality, e.g. pure and impure aesthetic or prudential deference etc.

²⁵ In a different paper, McGrath (2011, p. 114) widens the net of impure moral deference. She includes cases where the information transmitted is moral, but the deferrer's judgement is impaired in some way and they share the moral sensibility of the testifier; or cases where a moral dispute is solved by asking an unbiased stranger. She deems such cases "as straightforward as non-moral deference" (McGrath, 2011, p. 114). I do not agree with this verdict. At best, it is not *as* problematic as, e.g., deference about fundamental values. We will see why when I develop my positive argument

²⁶ Some philosophers restrict their discussion to thin moral concepts because they think it complicates the discussion (e.g. Fletcher, 2016, p. 48; Lord, 2018, p. 73). I will not do that because the account I develop can handle the complications. See discussion below.

to defer to one's doctor about what medicine to take but, by contrast, it seems odd to defer to one's ethics teacher about the moral permissibility of abortion. Since the aim of the debate is to find out whether the negative intuitions about moral deference are warranted (McGrath, 2009, p. 323), it seems more fruitful to look specifically at moral rather than non-moral deference, in order to isolate the features that seem to cause trouble.

However, there is a noteworthy question here about whether we can always draw a clear line between what counts as moral and what counts as non-moral information. The lack of a clear line transpires especially when we think about how moral deliberation occurs in the wild. In actuality, it seems that sometimes, from the first-person perspective, we cannot be certain of the borders between the moral and the non-moral. We just need to decide whether to have an abortion, whether joining the protest is the courageous thing to do, whether to take the job promotion, or whether it is bad to lie to our friend. Many moral and non-moral aspects come into play and it will be difficult to consciously separate them. Moreover, we might also not care about separating them, since all we want is to be able to make a decision.

Although this is true, the important bit for our purposes is whether we can, *in principle*, separate them. There will be instances where that will be fairly clear-cut, but not all cases will be like that. For example, the fact that animals feel pain is relevant to whether or not factory farming is morally permissible. That animals feel pain is clearly a piece of empirical information. So, if we are reflecting on the moral permissibility of eating meat, we might be able to separate the moral from the non-moral aspects of the problem in this way. But take another example, from Jones (1999). That someone will be upset by a certain action is, in a sense, a non-moral, empirical matter that might contribute to whether that action is right or wrong. But it might also be considered an *evaluative* matter:

Being upset contrasts with what is colloquially called "having a hissy fit," since the latter implies that the distress is *prima facie* unjustified. 'Upset' may thus have connotations of *prima facie* justified (or at any rate not *prima facie* unjustified) distress. (Jones, 1999, p. 61, fn. 11).

A similar problem might arise when we are dealing with thick moral concepts, e.g. fair, courageous, cruel, where we have both descriptive and evaluative aspects. So perhaps there are cases where the moral elements cannot, even in principle, be separated from the non-moral elements. Call cases of moral deference about such intertwined moral and non-moral matters *mixed moral deference*.²⁷

The existence of such cases may purportedly be a problem for pessimism. As Jones (1999, p. 61, fn. 11) notes, at least some pessimists are committed to saying that only pure moral deference is problematic. Because of that they will have to hold that the line between moral and non-moral information can be clearly drawn. If they deny this, they are unlikely to be able to give verdicts regarding cases where moral and non-moral aspects are intertwined. However, I am not sure this is true. That there can be tricky, borderline cases where the moral will be — in practice, but perhaps even in principle — inseparable from the non-moral, seems true, and pessimists need not deny this. Accepting this is not fatal, as they have the resources to respond to this challenge and to still appropriately diagnose different cases. For example, depending on how morally charged a case is, there might be pessimist views which will say that different instances of deference will be more problematic than others, i.e. allow for degrees (e.g. McGrath, 2009, p. 322). Then deference about these mixed aspects will just count as more problematic than deference about purely non-moral information, but less problematic than pure moral deference.

I am not too concerned with this *prima facie* problem because my version of pessimism will be able to accommodate it. As I will show later, my claim is that the relevant contrast is not between moral and non-moral deference but between deference which interferes with the capacity for practical deliberation and deference which does not. As such, I do not need to hold that there is a sharp line between the moral and the non-moral. However, I will follow the literature and start by focusing on *pure* moral deference just so that we can get to the heart of the problem. After we will have identified why moral deference triggers negative intuitions, we will be in a better position to say something about mixed moral deference as well, which I will do in chapter IV. From now on,

²⁷ I introduced and used this term in my MA thesis (Covaci, 2015, p. 16).

unless stated otherwise, whenever I discuss moral deference I mean pure moral deference.

2.2. The Right Kind of Bad Deference and The Wrong Kind of Bad Deference

Another point of agreement between pessimists and optimists is that some instances of moral deference are indeed problematic. The contention is, however, whether or not that shows us something about moral deference taken on its own, such that it would be recommended, in some sense to be specified, that we avoid it or not. Moral deference may seem problematic at first sight, but we have to be careful with the instances where it is not the fact that someone has deferred that triggers our suspicion, but where some other aspects of the case do that (Wiland, 2017, p. 57). Blanchard (2019, p. 1143) calls the adequate target for the pessimism *moral deference per se*, and explains the difference in this way:

Cases of moral deference are frequently plagued by features that make deference inappropriate but do not derive from the distinctly *moral* features of the case. For example, MORAL STATUS involves deference about a moral matter that ought to be obvious, and it may be inappropriate to defer about something obvious. Likewise, some cases of moral deference (for example, deference about whether it is morally permissible to eat meat, which is a popular example in the literature) are plagued by the problem of expert and peer disagreement. Perhaps it is inappropriate to defer about something subject to widespread expert and peer disagreement. However, obviousness and disagreement impede appropriate deference no matter the domain of belief. To get the intuitive force of the cases, you have to think that the moral content in particular contributes to the inappropriateness of deference.

Although I take the difference explained here to be very important, I am not fully satisfied with the explanation because it is not sufficiently fine-grained. The MORAL STATUS example mentioned above is about someone deferring about whether men and women have equal moral status. Such an instance of deference

has, obviously, moral content. So, in this case it *might* be that MORAL STATUS is problematic because it has moral content. But it is also deference about something obvious, a basic moral truth of which one should be aware. So how can we know that this case triggers negative intuitions because of the obviousness of the content rather than the fact that the content is moral? The problematic character of this case is then either indeterminate or overdetermined. Referring to the moral content of some act of deference is not sufficient. Moreover, this issue can generalise to other cases of deference that are intuitively problematic, for example aesthetic deference. There, of course, there is no *moral* content that contributes to the inappropriateness of aesthetic deference, but an aesthetic one. So, this criterion to distinguish between the kind of deference that we want to focus on versus the one that misses the target would not work there.

Nevertheless, the distinction between these different cases of problematic moral deference discussed in the quote above is crucial, because if we are not discussing the same types of examples, then we are just talking past each other — which is precisely what happens in the literature sometimes, with respect to this issue, as I will show in chapter III. Optimists think that we only have cases of moral deference that are problematic because of some pre-existing feature of the agent (e.g. a character flaw) or of the topic of deference (e.g. that the topic is controversial). They contend that once we explain away these cases, we see that moral deference is not actually bad. Pessimists, however, although they can accept the existence of such cases, (should) think that they do not exhaust the range of problematic cases. Not all problematic cases of moral deference are problematic merely because they uncover some pre-existing issues with the deferrer or the topic. There are instances of problematic deference that trigger negative intuitions because of what they cause or contribute to,²⁸ taken on their own. We just need to find the right cases to illustrate it.

²⁸ Different views will spell this out differently, e.g. deference is problematic in virtue of affecting the agent, the belief, or the action formed on its basis. Fundamentally, no one suggests that the very act of deferring as such is problematic in itself, but only because of its alleged bad consequences. In this sense, deference is not problematic *qua* deference. Nevertheless, the point raised here still stands and I think all of this is intuitive enough to see that there is a distinction between the problematic cases optimists discuss and those on which pessimists (should) want to focus.

I think this is the crucial distinction here: regardless of content, we need to see whether deference, on its own, causes or contributes to anything bad, and not whether it just reveals something bad about the agent who defers. I will thus refer to the former, after Blanchard, as cases whose problematic character tells us something about (moral) deference *per se*. This label is used strictly for simplicity and it should be read in light of what has been said here. It is not entirely accurate because I do not mean that moral deference is problematic *qua* moral deference, *intrinsically*, *as such*, or *in virtue of* being moral deference — that would entail that every single instance of moral deference is problematic. But if there are instances in which moral deference is problematic only because it reveals something about the agent (or indeed cases where moral deference is not problematic at all, in any sense) then moral deference will not be intrinsically problematic. But instead of having to tell this story every time I refer to these cases of problematic moral deference I will just say I am talking about moral deference *per se* or *in itself*— this is meant to exclude cases where moral deference is problematic because of some features that do not have anything to do with the damages to which deference contributes or causes (e.g. pre-existing character flaws).²⁹ Problematic moral deference *per se* is the right kind of bad deference on which the debate should focus. In what follows the discussion is about moral deference *per se* or *in itself*, unless stated otherwise.

2.3. There Is Something about Moral Deference

What do pessimists mean when they say that moral deference is “problematic” (McGrath, 2009, p. 322), “illegitimate” (Hopkins, 2007, p. 617), or “fishy” (Enoch, 2014, p. 229)? What do optimists mean when they say that there is no “general problem” (Sliwa, 2012, p. 177) with moral deference?

Prima facie, they make reference to some data or evidence that we have: certain instances of moral deference trigger or do not trigger some negative intuitions. Look at any pessimist view and you will find examples of moral deference that make us suspicious (take the aforementioned Eleanor case from

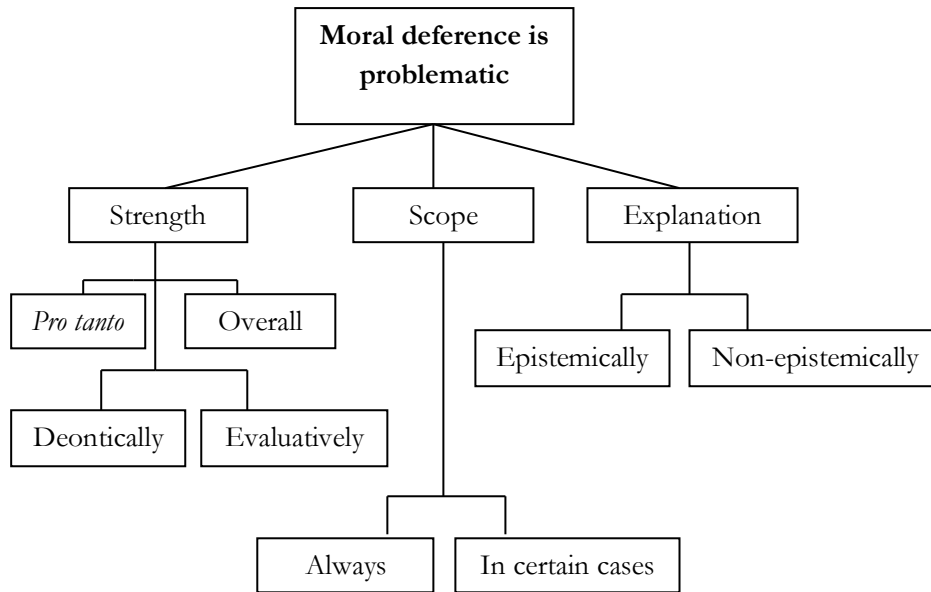
²⁹ Of course, moral deference can be bad even for those who already have some pre-existing flaws that trigger our intuitions when they defer. But in such situations the origin of our intuitions will be too unclear (or overdetermined), so I would say it is more fruitful to not focus on them.

Hills). But look at an optimist view and you will also find examples of moral deference that do not seem to elicit any negative attitudes from us (take Sliwa's (2012, p. 177) wedding example).

This data is not, however, unimportant. As I have said, the debate starts precisely with the perceived negative intuitions that arise around moral deference. And the current debate seems to be highly concerned with how the various views explain those intuitions, how they test against different cases of moral deference, and whether the verdicts and predictions they give match with our intuitions. A lot seems to hang on these intuitions that various cases of deference trigger. But there is a question about whether, and how much, it *should* matter that our explanations match with our intuitions. Are our intuitions informed and educated enough to be reliable? Do they track the distinctions that we will have to make regarding deference? Our explanations can, of course, be true even if they do not align with our intuitions. I will come back to this issue later, but for now bear in mind that one of the goals of pessimists and optimists is to try to capture the intuitions that moral deference triggers in us.³⁰

But pessimists and optimists also need to tell us what generates our intuitions and see whether these intuitions are warranted. That is, they must spell out what it means that moral deference is problematic or not. And I think that there are three main questions to be asked, or three main factors at play, here, which I will be examining in the following sections. One: how problematic or unproblematic is moral deference? That is, the dimension of *strength*. Two: when is moral deference problematic or unproblematic? That is, the dimension of *scope*. Three: why is moral deference problematic or unproblematic? That is, the dimension of *explanation*. Getting some clarity on these issues will take us closer to completing the task of defining moral deference pessimism and optimism. To help visualise all of this, the grid below illustrates the different moving parts of the pessimist claim that moral deference is problematic. They can, of course, be combined in many ways. The optimist claim that moral deference is unproblematic will just be the negation of whatever combination of elements a certain pessimist view holds.

³⁰ When some explanations do capture our intuitions, we can adjudicate which is the best one by seeing the extent to which it captures our intuitions, how well it does against different test cases, its general plausibility, the soundness of the argument on which it is based, and whether it can coherently extend to other kinds of deference to provide a unified account.



2.3.1. How Problematic (or Unproblematic) Is Moral Deference?

As I said in the beginning, roughly, pessimism has a negative attitude towards moral deference and optimism has a positive one — but what does this mean? We will see later how we can interpret what ‘problematic’ or ‘unproblematic’ mean, but until then we should examine the strength of the negative or positive claim we want to put forward. The strength dimension can be analysed along two further categories: one, we have the *pro tanto* versus the overall or all-things-considered category; two, we have the deontic versus the evaluative category.

The first category is pretty clear. One option is to say that moral deference is *pro tanto* problematic or unproblematic. That is, it can be problematic or unproblematic to a certain extent, with respect to whatever we think makes it problematic. This does not entail a blanket ban or permission on moral deference. Rather, it tells us that although we may have some reason to defer or not to defer — a *pro tanto* reason —, it can be overridden if weightier factors are at play. But, by definition, the *pro tanto* reason does not disappear, even if it is outweighed. By contrast, our claims about moral deference can also be all-things-considered or overall. This means that we consider all the relevant things that matter for the decision of whether to defer or not, and we see how they weigh

up against each other to come to a conclusion. Making an overall or all-things-considered choice is compatible with there being some *pro tanto* reasons that go against the final decision. For example, even if it might be *pro tanto* problematic to defer with respect to [insert your favourite pessimist argument here], it could still be all-things-considered unproblematic to defer. Or vice versa.

Then we have the second category. If we go on the deontic side, we could choose one of the following options regarding whether or not to defer: we ought; we have an obligation; it is morally required; or it is morally impermissible. If we take the evaluative side, then we could claim that it is good, bad, better, or worse to defer or to not defer. And, of course, we can combine this category with the previous one, such that we get a *pro tanto* or an all-things-considered obligation or impermissibility, or a *pro tanto* or all-things-considered goodness or badness (or ‘better than’, respectively ‘worse than’).

In the existing literature, very few people use this language or make it clear where exactly they stand regarding these categories. Crisp (2014) is the only one who explicitly opts for a *pro tanto* evaluative claim, but, alas, his view is not very developed. Howell (2014, p. 390) can also be plausibly interpreted as subscribing to a *pro tanto* evaluative view, albeit implicitly (his choice of word is *sub-optimal*). The same holds for Hills (2009, p. 98) as she says we have strong reasons not to take moral testimony. Lord (2018, p. 88) simply takes the pessimist claim to be a deontic one; more precisely, he interprets pessimism as saying that we generally (presumably *pro tanto*) ought not to defer. But he also argues that such an obligation-based claim is too strong and so he does not call himself a pessimist. Nevertheless, he seems to believe that moral deference is *pro tanto* bad — in his words: “there is something defective about deference” (Lord, 2018, p. 88) — to the extent that it does not get us acquainted with moral properties and precludes us from acquiring appreciate knowledge. Hopkins (2007) and Nickel (2001) discuss a *requirement* of morality that makes it such that we need to grasp the grounds of our moral beliefs. This suggests that they might prefer a deontic claim. Skarsaune (2016) focuses on all-things-considered claim, investigating the cases where the badness of deference outweighs all other aspects of the situation.³¹

³¹ Specifically, he claims (Skarsaune, 2015, pp. 354-356) that having authentic interactions with others can sometimes outweigh other values (he mentions the risk of wrongdoing, as he

On the optimist side, Enoch (2014) argues that in certain cases, specifically where there is a risk of wrongdoing to others, it is all-things-considered morally required to defer. Wiland (2017) and McShane (2018) appear to take the *pro tanto* evaluative path, proposing that moral deference brings about certain moral goods that cannot be otherwise obtained, such that it is good, in this respect, to defer. Overall, the general sentiment on the optimist side seems to be that moral deference *per se* is unproblematic *simpliciter* (i.e. neither specifically all-things-considered nor specifically *pro tanto*) — although this is not explicitly addressed by the optimists at any point.³²

So we can see that there are different options available, in both the pessimist and the optimist camp, but we can already make a few observations about the categories discussed here, and eliminate some candidates. No view quite hits the extreme. To say that moral deference is morally impermissible or that we have an obligation, be it *pro tanto* or overall, to not defer seems too strong.³³ We will see when we look at the scope of the claim, in the next section, that everyone involved in the debate allows for exceptions. So, it cannot be the case that moral deference is impermissible. The same holds for an overall obligation. A reasonable version of pessimism will have to admit that there are instances where the best thing to do just is to defer. The value of non-deference is outweighable. As for the idea that we have a *pro tanto* obligation to not defer, it is unclear what would justify it. Why have a *pro tanto deontic* rather than an *evaluative* claim? It is still unclear to me why Lord (2018) interprets pessimism in deontic terms.³⁴ The same can be said about optimism. To say that moral deference is morally required, without restrictions, would be implausible. There might be specific instances where it is the case that it is required to defer, but no one says, or should want to say, that this just applies generally. It is unclear what

responds to Enoch) in cases where we interact with people close to us and when the circumstances hold low stakes.

³² We could also say that the optimists think that certain cases of moral deference are *prima facie* problematic. This means that they seem problematic at first sight, but actually they are not; at least not *per se* because they can be explained away. See section 2.2 above.

³³ Although Nickel (2001, p. 256) sometimes uses strong language: “it is sometimes morally wrong for her to allow herself to get in a position in which dependence on moral testimony is necessary. For this reason, she may be morally culpable if the testimony leads her astray.”

³⁴ Perhaps it is considered that a *pro tanto* evaluative pessimism is not interesting (cf. Reisner and Van Weelden, 2015)? I discuss this in chapter III.

would warrant such a claim too.³⁵ So I would say that defining optimism and pessimism along an evaluative dimension is best. But whether to do so in a *pro tanto* or *all-things-considered* way is intertwined with questions about scope — more on this in the next section.

2.3.2. When Is Moral Deference Problematic (or Unproblematic)?

When is moral deference problematic or unproblematic? If moral deference is in some way problematic, or perhaps unproblematic, then the question of scope arises: do we want to say that all instances are problematic, all instances are unproblematic, or only some are problematic while others are not? In other words, we can choose between claiming that moral deference *per se* is always problematic, never problematic, or sometimes problematic (in some respect, in a *pro tanto* or all-things-considered way, evaluatively or deontically).

We can start by eliminating the extreme claims again. No reasonable pessimist would want to hold that moral deference *per se* is always all-things-considered problematic no matter what.³⁶ Likewise, no reasonable optimist should want to say that moral deference *per se* is never, to any degree, problematic no matter what.³⁷ In fact, the optimist needs only to deny whatever claim the pessimist holds. In the literature, it seems that there are at least a couple of better options to consider. One claim is that moral deference is always *pro tanto* problematic. Another one is that moral deference is always *pro tanto* problematic *in certain cases* (or: sometimes *pro tanto* problematic).

It is not always very clear where the current pessimist views stand. It is typically said that moral deference is acceptable when *unavoidable* (Hopkins, 2007, p. 617; Mogensen, 2015, p. 263). That is to say, when the person cannot, for various reasons, figure things out on their own. Pessimists agree with the

³⁵ No optimist supports this, but one particular view does get close to it, when it comes to certain situations: namely, Raz's position on deference, taken from the debate on legitimate authority. I will discuss this at length in chapter III.

³⁶ Mogensen (2015, p. 263) goes on to say that "pessimists don't believe that it is always inappropriate to defer in deciding moral questions. Nor need they think that it is typically inappropriate." I am not sure what he means by the second claim and I do not think the rest of the pessimists would subscribe to that. But the first claim is characteristic of pessimism (albeit still ambiguous between *pro tanto* and all-things-considered).

³⁷ Although they are quite unclear on this such that it often seems that they would want to hold this claim. See chapter III.

optimists that it is not problematic for those who do not have the capacity or the moral competence to make decisions on their own to defer (Mogensen, 2015, p. 264). Psychopaths and sociopaths would also better defer (Howell, 2014, p. 390). People for whom moral knowledge or moral understanding is decidedly out of reach can also unproblematically defer (Hills, 2009, pp. 123-124). The same applies to those in a situation of moral uncertainty (Mogensen, 2015, p. 263), especially if the situation is such that there will be grave consequences if a wrong decision is made (Howell, 2014, p. 390). As Howell (2014, p. 390) puts it:

if one is in a position to learn a moral fact by deferring, and one cannot come to know that fact non-deferentially without substantial cost, it might well be that almost always one should do so. Otherwise, one is apt to do something morally impermissible.

So the main point is along the lines of ‘ought implies can’. It is fine to defer “if we know that we’re unable to resolve a difficult moral issue and reasonably expect that others can do better” Mogensen (2015, p. 264).³⁸

However, it is unclear whether the existing pessimists want to say that although all-things-considered unproblematic, moral deference still is *pro tanto* problematic in these unavoidable cases. For example, Hopkins (2007, p. 617) says: “Pessimism is defined by the claim that it is illegitimate to rely on the word of others in moral matters (unless such reliance is unavoidable)”. But he does not make it clear whether those cases where deference is unavoidable still count as problematic in any sense. Hills’ (2009, pp. 123-124) view on this is also ambiguous:

If moral understanding is unavailable for you, there is no point in your trying to acquire and use it. Given your situation, morally worthy action and proper orientation with regard to your moral beliefs are out of reach. But if you are lucky and you have access to reliable people whom you can trust, you might be able get moral knowledge and, as a result, do the right thing. Since doing the right thing is very important, you should trust moral testimony from trustworthy and reliable sources

³⁸ Is this true even in situations where we do not have to necessarily make up our mind, e.g. in a non-urgent situation; or regarding a speculative question? Some pessimists would deny it. Hills (2009, p. 124), for example, claims that if there is a chance that we could gain moral understanding we should try and do that rather than deferring.

and defer to moral experts if you cannot gain moral understanding.

Howell (2014, p. 390) is the only one who seems to strongly suggest that when it comes to the conceded cases, moral deference is all-things-considered unproblematic, but nevertheless *pro tanto* problematic:

An asymmetry is preserved just so long as deference is somehow sub-optimal to the extent that it would have been better if the knowledge had been attained without deference. It could be that this value of non-deference is almost always outweighed by the risk of performing an impermissible act, but that value remains.

In other words, this would be an ‘always *pro tanto* bad’ claim about moral deference on the pessimist side because although there are concessions, the *pro tanto* badness never goes away. It is not clear to me whether all the existing pessimists would subscribe to this.

The other option is to say that in the unavoidable cases moral deference is actually unproblematic *simpliciter*. That is, not problematic in any sense, not even a *pro tanto* one. The motivation for this might be that such unavoidable cases seem to not trigger negative intuitions. Take this case from Sliwa (2012, p. 178):

TRIP: Anna is a journalist who is preparing to go on a reporting trip to a dangerous and conflict-ridden area. She has to tell her family that she will be away but she really doesn’t know what to tell them. If she tells them where exactly she’s going and why, they will be extremely worried. On the other hand, she worries that by evading the questions she would be lying. She goes back and forth but cannot decide what the right thing to do is and eventually decides to ask a friend whose judgement she trusts.

This is an unavoidable case, as Anna is in a situation of moral uncertainty. Sliwa diagnoses this as an unproblematic case of deference — it does not seem to trigger negative intuitions. Definitely not in the way Hills’ (2009) Eleanor does. So perhaps it would be at least a little odd for a pessimist to say that this case is nevertheless problematic, albeit only in a *pro tanto* sense. So, on the one hand, if

we feel the pull of having to match our pessimist explanations to our intuitions, such a verdict might not sit well. On the other hand, we could also just say that our intuitions are indiscriminate between *pro tanto* and all-things-considered badness. Just because it *seems* to us that it is fine for Anna to defer, it does not mean that there is indeed nothing even *pro tanto* problematic about it. It is too difficult, if not impossible, to say what exactly our intuitions are responding to there (if it is widespread that this instance of deference does not trigger negative intuitions).

Nevertheless, if someone is worried about this,³⁹ then they could simply deem the unavoidable instances unproblematic *simpliciter*, and choose to focus on another set of cases: those where deference is avoidable. That is, cases where it is possible for people “to settle moral questions for themselves or to remain agnostic” (Hopkins, 2007, p. 613). The claim would then be the following: it is always *pro tanto* problematic to defer, *in certain cases*, namely, cases where one could avoid deferring.

On the optimist side, things are not a whole lot clearer. As their main goal is to defend moral deference, many optimists do not mention whether they would concede that at least sometimes moral deference is problematic in itself (Sliwa, 2012; Groll and Decker, 2014).⁴⁰ Nevertheless, some do explicitly say that moral deference is at least regrettable (Enoch, 2014; Jones, 1999, and Lord, 2018, if you class him as an optimist). Others even say that moral deference is sometimes distinctively *pro tanto* good (McShane, 2018; Wiland, 2017). But since optimism is mainly the denial of pessimism I will later reconstruct it as such.

2.3.3. Why Is Moral Deference Problematic (or Unproblematic)?

Finally, when it comes to our claims about moral deference, be them negative or positive, we have different explanation for what makes it problematic or unproblematic. Start with the negative side. Pessimism has two main strands. *The epistemic strand* takes moral deference to be epistemically problematic and *the*

³⁹ One need not be worried about this, or they can try and argue against the *prima facie* intuitions and say that actually deference in such cases is nevertheless *pro tanto* problematic. I am only presenting the possible options here.

⁴⁰ Although they do not deny moral deference can be problematic, but just not *per se*, see section 2.2. above.

non-epistemic strand takes it to be non-epistemically, typically morally or evaluatively, problematic. Although this thesis is not concerned with the epistemic strand, I will give a quick overview of the issues it raises so that the picture of the debate is complete. After that, I will discuss the non-epistemic strand.

The epistemic strand of moral deference pessimism finds reasons to object to moral deference based on various issues surrounding the notion of moral expertise.⁴¹ It is at least implied that we should not take moral testimony because it is epistemically problematic. The two general views are the following: 1) there are no moral experts; 2) moral experts might exist, but we cannot reliably identify them or recognise whether they meet adequate epistemic conditions for the transmission of moral testimony (e.g. being deserving of trust). The first type of view is less common than the second, and it usually goes along metaethical lines.⁴² One could be an error theorist and think all our moral judgements are false. One could be a certain sort of non-cognitivist and think that there is nothing to transmit via testimony since moral judgements express non-cognitive states of the speakers. Philosophers discussing such views cite Ryle, who thought that morality is a matter of caring rather than knowing, so there can be no expertise in this area (Driver, 2013; Singer, 1972).

The second type of view posits that moral experts do exist, but that they are difficult to identify or enter in a relation of trust with. For one, it is not as easy to evaluate the credentials and track records of a moral expert, like it is with a doctor or an accountant. With the latter, you have hard evidence of their previous work, you can see how their judgements and actions turned out, verify the accuracy of their predictions, and so on. We cannot do the same with moral experts as we do not have a similar objective track record (Hoffmann, 2012). Moreover, how can we, as non-experts, judge the track record of the alleged expert since we are not experts ourselves, which is precisely why we need help? (McGrath, 2009, pp. 334-335). It seems that there is no safe way for a non-expert to identify a moral expert, even if the latter does exist. Plus, there is the further

⁴¹ A different epistemic argument against moral deference comes from Simion (2018). She argues (2018, p. 483) that pure deference is bad because it violates a norm of assertion: "One's moral assertion that p is epistemically permissible only if (1) one knows that p and (2) one's assertion is accompanied by an appropriate explanation why p".

⁴² I discuss the possible metaethical explanations for the problematic character of moral deference in chapter VI.

question of whether the moral experts will be reliable in transmitting the right moral judgement to us. We need to know the adviser is reliable in delivering the right advice and not just that they are able to figure out the answer. But it is difficult, it is argued, to recognise whether they are being impartial and are putting their self-interest aside (Driver, 2006, pp. 631-632). Moreover, when it comes to moral testimony, the standards for trusting someone are particularly high (Jones, 1999, pp. 72-74).

I think various successful defences against both these types of arguments have been put forward (e.g. Hopkins, 2007; Howell, 2014; Lord, 2018). I agree with Hills (2009, pp. 96-97) that these epistemic issues are not insurmountable:

Few people, if any, can properly be described as experts on all moral matters. But to discover someone who has better judgement or more experience than you about some single moral issue and who seems generally trustworthy is not that difficult.

I choose not to engage with these arguments myself because this thesis is, by stipulation, about the non-epistemic status of moral deference. Even if we want to be sceptical of the existence or the possibility of identifying moral experts, we can still ask the central question of this thesis in a conditional form: *if* there are moral experts and there is a way to reliably recognise them, are there any non-epistemic reasons not to defer to them?⁴³

So this leaves us with the non-epistemic strand of moral deference pessimism. This type of view holds that moral deference is morally, normatively, or axiologically problematic, in a *pro tanto* or all-things-considered sense, always or sometimes. A few options are on the table. The problematic character of moral deference can be explained in terms of a failure to give us moral understanding (Hills, 2009). Since moral understanding is instrumentally valuable, as it is necessary for moral worth, virtue, and being able to justify ourselves to others, it is bad to defer about moral matters. Moral deference has also been said to be defective because it precludes us from obtaining complete virtue (Howell, 2014) or performing morally worthy actions (Nickel, 2001).

⁴³ I will simply be assuming that all the epistemological stars align, and moral testimony can in principle work epistemically. This implies further assumptions, e.g. that moral knowledge exists and that we can transmit it via testimony. If any reader disagrees with them, they should plug in some conditional here as well.

Another argument is that moral deference undermines the important value of authenticity (Mogensen, 2015; Skarsaune, 2016). Or it does not give us appreciative knowledge, a type of knowledge central to our moral lives (Lord, 2018). These are the main ways the pessimists explain the ‘problematic’ character of moral deference, which I will discuss at length in the following chapter.

Optimism denies all of this. More precisely, various versions of optimism deny all the various pessimist explanations that purport to explain what it means that moral deference is problematic — most plausibly interpreted as being about moral deference *per se*, as explained above. There are three main ways in which they go about doing this. One is by arguing against specific pessimist arguments (e.g. Sliwa, 2012; McShane, 2018), and say they do not work. The second is by showing that some cases of moral deference that trigger negative intuitions do so because of some features of the agent and not of moral deference (Groll and Decker 2014; Sliwa 2012); that is, although some instance of moral deference can indeed trigger negative intuitions this does not mean that moral deference *per se* is problematic, as discussed above. And, finally, by illustrating that moral deference need not always trigger bad intuitions, and that it brings distinctive moral goods, is better than non-moral deference (McShane, 2018; Wiland, 2017), or is sometimes even morally required (Enoch, 2014). I discuss and engage with the optimist strategies at length in chapter III. This is just a preview into the main ways in which the ‘why’ of the (un)problematic character of moral deference has been explained.

3. Moral Deference Pessimism and Moral Deference Optimism

Having examined the variables that can change from view to view, I think we can extrapolate some fairly clear minimal definitions of moral deference pessimism and optimism. I take the descriptions that I will offer to follow the spirit of each side, even though perhaps not their letter exactly.

So, for now, the best interpretation of the existing moral deference pessimist stance can be formulated in the following way:⁴⁴

⁴⁴ This is based on Crisp’s (2014, p. 141) formulation of pessimism: “in certain cases, judging on the basis of testimony is always worse, to some degree, than judging for oneself, even though in those very cases, and indeed others, it may be better overall to rely on testimony”.

Moral Deference Pessimism	Moral Deference Optimism
For any case C in a certain range R (to be specified), moral deference is always <i>pro tanto</i> bad in C, even though it may be all-things-considered better to defer in C (and in many other cases, too).	It is not the case that for any case C in a certain range R (to be specified), moral deference is always <i>pro tanto</i> bad in C, even though it may be all-things-considered better to defer in C (and in many other cases, too).

The certain range R is left open for pessimists to choose: it could be limited to cases where the deferrer could have made the judgement on their own, or expanded to cases where they could not have done so. I believe that this is the minimum that the pessimist needs in order to be an interesting position, but that it is also the best interpretation of the current literature.⁴⁵

Typically, the pessimist views have as a corollary *the asymmetry thesis* (originally developed in Hopkins (2007), label from Groll and Decker (2014)), according to which there is a contrast between different types of deference, such that some are problematic and others are not (or that some are more problematic than others). Think of the examples I started with. It does not seem problematic to form beliefs about how to invest our money based on the testimony of an accountant, but it does seem problematic to form beliefs about the moral permissibility of capital punishment by deferring to our reliably virtuous friend. Or, as we have seen above, it seems fine to rely on Google Maps, while it would not seem fine to rely on its moral counterpart, Google Morals. Most optimists also subscribe to the denial of the asymmetry thesis:⁴⁶ they hold that there is no contrast between different types of deference, such that some are problematic

⁴⁵ There might be another way of conceptualising pessimism. We can say that moral deference is always *pro tanto worse* than not deferring or than judging for oneself (as Crisp formulates it). But a noteworthy aspect of this claim is that, on its own, it does not entail that moral deference is *pro tanto bad*. This may not be sufficient, then, to capture the negative intuitions that moral deference triggers. Recall the Google Morals example. The intuition does not seem to be merely that it is *worse* that the person uses Google Morals rather than judging for themselves. Instead, “there seems to be something wrong with using Google Morals” (Howell, 2014, p. 389). The ‘worse than’ claim also contrasts with the general language used in the debate, where it is said that moral deference is fishy or inappropriate, which suggests that it is, in some sense, bad (and not just worse than not deferring). For these reasons, I take the best version of pessimism to be based on a *pro tanto* badness claim.

⁴⁶ Except for Enoch (2014), Lord (2018), and also Wiland (2017) but in a different sense. I expand on their views later.

and others are not, or that some are more problematic than others. Moral deference and other types of deference are alike. And this is debate in broad brushstrokes. Now we can go into the weeds.

4. Concluding Remarks

The purpose of this chapter has been to give a general overview of the state of the debate on moral deference, to pave the way for constructing my own view. I hope to have shown that the existing landscape is not very neat. This warrants what I will do in chapter II, namely more reconstruction work such that I can engage with the best versions of moral deference pessimism. The aim is to show that, although the existing work has been insightful, it has not got to the heart of the problem of moral deference. This is what motivates and gives way for my own, novel, version of moral deference pessimism.

Chapter II: Foes of Moral Deference

In this chapter I argue that existing pessimist views about moral deference have failed to give a fully adequate explanation of the problematic character of deference. This will further motivate the need for a novel position that can properly achieve that task. I will discuss and reject the views that I take to be the most important and the most developed in the literature. These are what I call *the moral understanding strategy*, *the virtue-based approach*, *the acquaintance view*, and *the authenticity argument*.

In the previous chapter, I have argued that the best interpretation of moral deference pessimism is the following: for any case C in a certain range R (to be specified), moral deference is always *pro tanto* bad in C, even though it may be all-things-considered better to defer in C (and in many other cases, too). The corollary claim that is typically also held by pessimists is the so-called *asymmetry thesis*, according to which there is a contrast between different types of deference, such that some are problematic and others are not (or that some are more problematic than others). The existing versions of pessimism vary on how they specify what ‘problematic’ means and in what way moral deference is problematic, so their negative claims about moral deference differ in strength and scope. They also have different stances on whether they take moral deference to be uniquely problematic or just more problematic than other kinds of deference. But we can continue by taking the views to at least subscribe to pessimism as described above. In what follows I will start by presenting a general objection to moral deference pessimism, which applies to most, if not all, of the existing versions. This criticises them on the basis of their insufficiently developed and unclear commitments as well as because of the kinds of examples of deference they target. Subsequently, putting those problems to a side and assuming the best pessimist views, I move on to discuss why they are nevertheless unsuccessful.

1. A General Problem

A general problem for the existing pessimist literature stems from the fact that it is not always very clear about its commitments. For one, their stance on the kind of instances of deference that is worth investigating is not fully motivated and, in the end, it also sets the debate on the wrong path. Second, because their negative claims about moral deference are ambiguous, it is difficult to see how wide the net they are casting on moral deference is. As such, it is not easy to see how these views test against different cases of deference. All of these might not seem like *prima facie* fatal flaws, but I will show that the consequences have been substantial enough to merit discussion.

To start, see a case that is taken in the literature to be paradigmatic of moral deference:

Danielle hears about an upcoming demonstration protesting Israel's war in Gaza. Although she knows the causes of the war and knows that civilians are dying from IDF bombing, Danielle is unsure whether the war is just. She doesn't try to think through the matter for herself. Instead, she asks a reliable and trustworthy friend, who says the war is immoral. Danielle accepts her friend's claim and joins the protest. Asked by a journalist why she is demonstrating, Danielle says she knows the war is wrong because her friend told her so. (Mogensen, 2015, p. 261)

This is a typical example of moral deference that the current literature uses (see similar examples in Hills, 2009; Howell, 2014; Nickel, 2001; Sliwa, 2012) to generally fuel the debate. It involves one person deferring to another one, about one moral issue, at one time. I will refer to such cases of deference as *isolated*, *individual*, or *one-off* cases of deference. Isolated instances of deference are effective in triggering strong intuitions about moral deference. Even most of the optimists would agree that there is something problematic about them. But these cases do not do the work that the pessimists want them to do.

For one, such examples are usually under-described, which ends up causing trouble. We know about Danielle that she does not try, by herself, to reflect on whether the war is just. The same goes for other instances used in the literature (Hills, 2009, p. 94; McGrath 2009, p. 321). That is all we know. But I

think it is important to inquire into *why* the protagonists defer. This is because their reasons and motives might play a role in how we evaluate them as moral agents, which might show us something about what the source of our intuitions really is (this is what I was anticipating in chapter I, section 2.2).

Filling in the details of Danielle's case, one way to interpret her motivation to defer is to say that she is a lazy person, who does not like to think about complex issues and prefers ready-made answers. Perhaps she is also indifferent and thinks she can use other people to solve her moral conundrums. Simply put, we can imagine that Danielle has a flawed moral character. However, if this is the case, then we might have an explanation of our negative intuitions about her case that does *not* focus on the fact that Danielle has deferred. Rather, her deference reveals some moral shortcomings that she has had previously: she is morally irresponsible, or she uses her friend merely as a means, and this what makes us wary of her deference. Interpreted like this, this instance of moral deference does not seem problematic *per se*. It only uncovers a pre-existing problem with the agent, and that is what fuels our negative intuitions about the case.

Take another case which, although not under-described itself, makes the previous point in a clearer manner:

Ron is an extremist, believing that killing a person is not generally immoral but that killing a fellow Jew is a grave sin. Ron would like to kill Tamara, but he refrains from doing so because he wants to do the right thing, and he knows (on the basis of his rabbi's testimony) that the right thing to do is to refrain from killing her. (Hills, 2009, p. 115)

This instance of moral deference seems odd indeed. But, as Sliwa (2012, pp. 184-185) notes, this is because of Ron's extreme moral ignorance. His moral sensitivity must be deeply compromised if he needs help in figuring out whether to refrain from killing someone. Not to mention the fact that he usually thinks that killing is not immoral in general! As (Sliwa 2012, p. 185) puts it:

what our intuitions are latching onto in these cases is not that the agent is resolving her moral ignorance by relying on testimony, but rather that she is morally

ignorant in the first place. Some instances of moral ignorance are just problematic in and of themselves.

The problematic character of this case of moral deference is then *explained away* as not being about moral deference in itself or *per se*. Rather it is about a moral flaw that the agent had from the start, before he even deferred. But, as I have shown in the previous chapter, the relevant kind of problematic moral deference is the one that is problematic in itself. What I said before about Danielle is just the application of this explanation to her case. We can imagine that Danielle, although not as deeply morally ignorant as Ron, has a flawed moral character. But then we are not looking at the right kind of bad deference.

However, this is not the only way that under-described examples of isolated moral deference can be interpreted. Say Danielle (or any other moral agent used in examples in the literature) actually defers because she is in genuine moral uncertainty and knows she cannot solve the problem on her own. Or perhaps she has a bad track record at making this sort of moral judgements in comparison to her friend. She is generally committed to social justice and thinks this issue is important, so she wants to know the right answer and she defers.⁴⁷ If we look at these instances in this way, is it still plausible that they are problematic? Intuitively, it *prima facie* does not seem so. When we think that this is the best Danielle can do, negative intuitions are not so easily triggered. This is how the optimists would describe the case (Driver, 2015; Enoch, 2014; Sliwa, 2012) and many, perhaps all, pessimists seem to agree. The latter group is silent on whether they think such a case is nevertheless *pro tanto* problematic, as these are the unavoidable cases that they concede as acceptable (although this is a viable option for them — see chapter I, section 2.3.2).

The issue is that these intuitions themselves are ambiguous as to what they are tracking: they could be tracking a sense of unproblematic *simpliciter* or an all-things-considered sense of unproblematic.⁴⁸ I am saying they are ambiguous mainly because no pessimist explicitly addresses what they make of this distinction, as it does not even seem to be on their radar. But perhaps other

⁴⁷ Some philosophers use precisely this type of examples (e.g. Enoch, 2014; Hopkins, 2007; Jones, 1999; Sliwa, 2012)

⁴⁸ There might be a third option: *pro tanto* unproblematic. However, since the pessimists readily concede these cases as acceptable, with no explicit qualification, I am inclined to think that is the least plausible interpretation. But what I say here would apply to a *pro tanto* unproblematic claim too.

people have very clear and unequivocal intuitions about this. Nevertheless, it is worth making it explicit that here pessimists have a choice to make. As I have mentioned in chapter I, they can opt for the former and say they restrict their negative claim to cases where moral deference is not unavoidable, to use Hopkins' term again.⁴⁹ But even then we should still be investigating how such deference might happen. More precisely, we should ask why would a person who could figure things out alone defer? There could be many reasons, but we need to know in order to judge the case and see whether negative intuitions arise there too, and how they can be explained. This is because one reason could be that the person has a flawed moral character and it turns out that this is not a case of deference that is problematic *per se*. And then we are back at filling in details and seeing how the case ultimately turns out, which might just give rise to the same problem again. But pessimists can also opt for the latter option and say that, while such isolated instances of deference — be them unavoidable or not — are *all-things-considered* unproblematic, they are nevertheless *pro tanto* problematic.⁵⁰ I do not think this works though. Here is why.

It seems to me that whether or not an isolated instance of deference, like Danielle's, is *pro tanto* problematic heavily depends on the explanation we give for why moral deference is problematic. For one, the intuitions are not helping us here. A lot of people, from both the optimist and the pessimist side, agree that cases like Danielle's seem unproblematic (but we do not know whether *simpliciter* or all-things-considered), if the details tell us that she is morally uncertain, in a situation of urgency etc (i.e. that she is in one of those unavoidable situations).⁵¹ So we have to *argue* for the claim that, although our intuitions suggest, at least *prima facie*, that this is unproblematic deference, it nevertheless is *pro tanto* problematic. That is fine. We need not hold our intuitions, as they are, on a pedestal. They can be uneducated or simply wrong. So we may have to adjust or give them up. The issue is that, whatever the argument, it is very implausible that an isolated or one-off instance of deference will be damaging in

⁴⁹ Textual evidence seems to support that at least Hopkins (2007, p. 617; p. 621) and Mogensen (2015, pp. 263-264) opt for this.

⁵⁰ At least Crisp (2014), Howell (2014), Lord (2018), and perhaps Skarsaune (2016) seem to opt for this.

⁵¹ Again, perhaps some people *do* have strong intuitions that such a case is *pro tanto* problematic but all-things-considered unproblematic. What I say here is based on the literature and how both the pessimists and the optimists discuss this issue.

the relevant way. As we will see, once I explore the pessimist accounts in more depth, philosophers have cited lack of moral understanding (which would affect moral worth), precluding virtue, lack of appreciative knowledge, or the absence of authenticity, as reasons why moral deference is bad. But none of these moral values or ideals require perfect compliance: to be virtuous, for example, it does not mean that every token of our moral decision-making must be fully compliant with virtue (Hursthouse and Pettigrove, 2018).⁵² Isolated deference, the protagonist of the paradigmatic examples in the debate, is in principle *not* incompatible with reaching those values or the ideals associated with those values.

Moreover, perfect compliance with moral ideals might not even be something that we should strive for (Wolf, 1982). This is because a moral saint would be so dedicated to being moral that their life would be barren: they would have no time for any non-moral goods. As Wolf (1982, p. 421) puts it:

For the moral virtues, given that they are, by hypothesis, *all* present in the same individual, and to an extreme degree, are apt to crowd out the nonmoral virtues, as well as many of the interests and personal characteristics that we generally think contribute to a healthy, well- rounded, richly developed character.

Thus, I think that when it comes to *isolated* cases of moral deference we cannot claim that they must be nevertheless *pro tanto* bad even if all-things-considered fine. For this reason, I think that the pessimists do not succeed in establishing the ‘always *pro tanto* bad’ claim with respect to individual cases of deference.

So, the individual cases of deference used in the literature do not seem to serve us very well. Now I do not deny that there might be some isolated instances of moral deference that are at least problematic *per se* and for which

⁵² There might be a worry here that this will not hold for the moral understanding view, where it is said that moral understanding is necessary for morally worthy actions. Every single action based on deference will lack moral understanding and so will lack moral worth on this view. So isolated cases can be *pro tanto* bad in that respect and to that extent. To this I say the following. First, as we will see later, it is not at all clear that every action based on deference lacks moral worth completely. Or, if we want to say that it lacks higher degrees of moral worth it is unclear whether that is something worth pursuing. Secondly, and more importantly, on the moral understanding view, performing morally worthy actions is part of a more comprehensive view on how one should be oriented towards morality and what makes one a good moral agent. But it is implausible that to be adequately morally oriented, and a good moral agent, one’s every single action needs to be (highly) morally worthy.

our best explanations can actually account. But this will be established when we investigate these cases closely as they, and the intuitions they trigger, will depend on the details of the case, particularly on *why* the agent defers. Although not necessarily *impossible*, it is surprisingly challenging to construct a detailed enough isolated case of deference that we would consider problematic and which triggers negative intuitions unequivocally only in virtue of being an instance of moral deference. But if we do find such cases, our explanations will also have to be such that they manage to cover isolated cases and not face issues like those I have discussed above. Until then I hold that we will not learn a lot about whether moral deference is indeed problematic *per se* by investigating isolated cases of deference. Because of all the issues they encounter, such cases are just not that useful for the debate. So I propose we leave them behind. This does not mean, however, that optimism has been established. Rather, I think we can take a different and novel approach and show that, while diagnosing individual instances of deference is, to say the least, unfruitful, we can focus instead on what I call *recurrent moral deference*, i.e. moral deference that occurs repeatedly in an agent's life.⁵³

2. A Way Forward: Recurrent Moral Deference

Not many philosophers use examples of recurrent moral deference (the only exceptions are Howell, 2014, McShane, 2018, and briefly Skarsaune, 2016). Those who do, nevertheless fail to differentiate it clearly from isolated deference and to affirm its distinctly problematic character (additionally, their cases also suffer from the under-description charge discussed above.) Take again Howell's (2014, p. 389) Google Morals scenario that I have introduced in the previous chapter:

Suppose those wizards at Google come out with a new app: Google Morals. No longer will we find ourselves lost in the moral metropolis. When faced with a moral

⁵³ Not all cases of recurrent moral deference will be adequate. For instance, negative intuitions regarding recurrent deference about basic moral truths, e.g. whether torturing others for fun is wrong, can still reflect character flaws (in which case the problematic character of deference is most likely overdetermined). The point is that there will be a wide range of cases where recurrent moral deference seems problematic, but not on this ground.

quandary or deep ethical question we can type a query and the answer comes forthwith. Next time I am weighing the value of a tasty steak against the disvalue of animal suffering, I'll know what to do. Never again will I be paralyzed by the prospect of pushing that fat man onto the trolley tracks to prevent five innocents from being killed. I'll just Google it. Again I find myself dependent on Google for my beliefs, but in this case it seems to many, myself included, that this is not a good way to go. There seems to be something wrong with using Google Morals.

Indeed, it does seem like there is something wrong with using Google Morals, particularly if done *recurrently*. To see the contrast, imagine he only uses the app once. Would it be just as problematic? Well, we are back where we started because we need to ask why he is deferring: is he lazy or irresponsible? Is he in genuine moral uncertainty and cannot figure it out alone? Moreover, these questions can be asked of the recurrent Googler, as the case is unspecified with respect to his motivations too. So I think we can do better by constructing even sharper examples that eliminate the need for such questions:

JOHN. Suppose John is nice and considerate, tries to do the right thing as often as possible, and he wants to be virtuous in an unfetishistic way, just as much as we all do. Most of us would be inclined to call him a good person. However, John defers quite a lot. He sometimes needs help with seeing which considerations are morally relevant in a situation. Other times he does not know how such reasons weigh up. Once in a while he just needs to simply adopt his reliable friend's moral judgement because he cannot make up his own mind and cannot understand his friend's explanations of the right answer.

If the description is found lacking, we can *stipulate* that John does not defer because his moral or epistemic character is acutely flawed, he is not deeply morally ignorant, and he is not to blame for his need to defer. He just finds himself in situations where he genuinely needs help. I think that, although John is the kind of person whom we would ordinarily consider good, something is not quite right about this pattern of moral deference that exists in his life. His recurrent deference is problematic in a way that him deferring only on one

occasion would not be.⁵⁴ Perhaps each occasion is not problematic at all on its own, yet his repetitive deference seems fishy. Such cases of recurrent moral deference, I hold, escape the problems that individual cases encounter and thus are the right target to focus on if we want to assess moral deference *per se*. Patterns of recurrent cases are more likely to interfere with our ideals, especially by comparison to individual cases (though this still does not mean that we should aim to reach those ideals, see Wolf, 1982). Moreover, this is not limited to the moral case. If we want to evaluate deference *per se*, we always must look at recurrent cases.

However, one might think that John just happens to be unlucky such that he needs to repetitively defer and, because of that, there is nothing bad about his recurrent deference. Although I agree that John may be unlucky, I do not think that this tells us anything about moral deference. Perhaps it tells us that John might not be blameworthy for his need to defer, depending on whether we think ignorance exculpates. But I do not see any reason to think that blameworthiness of agents who defer necessarily goes together with, or is conceptually related to, the problematic character of moral deference *per se*.⁵⁵ Even if one does not agree and does not share my negative intuitions about John, in chapter IV I argue that one should, i.e. that John's recurrent deference is indeed (*pro tanto*) problematic.

So I formulate a new and improved version of the (moral) deference pessimist claim. This applies to moral deference, but also wider, to any kind of deference:

⁵⁴ It is possible, however, that recurrent deference sometimes exposes a character flaw even *more* clearly. That is, the character flaw is precisely why one needs to defer repeatedly. This might be true (although in John's case it is stipulated that it is not). But not all cases will be like this. Additionally, I will argue that recurrent deference is always *pro tanto* bad to the extent that it interferes with the capacity for practical deliberation; so even if a character flaw is present, deference will further be problematic for this reason too (although it may be difficult to tell given the possible overdetermination of our intuitions).

⁵⁵ This challenge is similar to that which arises in the unavoidable isolated cases of moral deference: our *prima facie* intuition is that, since ought implies can, why would deference be problematic? That is, such deference seems all-things-considered unproblematic. The difference is that it is implausible to hold that an unavoidable isolated case of deference is nevertheless *pro tanto* problematic, for the reasons given in the previous section. However, it is not implausible that recurrent deference is nevertheless *pro tanto* problematic, even if unavoidable and so all-things-considered good, in virtue of being recurrent.

Deference Pessimism
Recurrent (moral) deference is always <i>pro tanto</i> bad, even if deferring is not <i>pro tanto</i> bad in individual cases, and even if some patterns of recurrent (moral) deference are all-things-considered better than not recurrently deferring in those ways.

The existing pessimist views have two options here. They can try to improve their isolated cases of moral deference and defend themselves against the challenges I posed here. Alternatively, they can embrace my version of pessimism and modify their targets and their arguments accordingly. For the reasons I presented above, I am not certain they can take the first option. And for reasons I will present below, I am not certain they can take the second option either simply because there are independent grounds for thinking that their arguments are not fully successful. In chapter IV I will argue for a new pessimist argument against moral deference that I hope will advance the debate.

But, for now, let me proceed by considering the existing pessimist positions on their own terms. That is, by interpreting pessimism as claiming that for any (isolated) case C in a certain range R (to be specified), moral deference is always *pro tanto* bad in C, even though it may be all-things-considered better to defer in C (and in many other cases, too).

Various arguments against moral deference have been proposed.⁵⁶ Although all of them touch upon interesting points, I will focus here only on the most promising positions. In what follows I will thus investigate what I take to be the four main strategies in the literature: *the moral understanding strategy*, *the virtue-*

⁵⁶ One fairly natural argument that I do not discuss at length is that moral deference is bad because it infringes upon our autonomy, as we give to others the power to make our decisions for us (Wolff, 1970). However, I take the objections that have been brought to this position to be convincing (B. Cross, 2017; Driver, 2006; Lillehammer, 2014). For example, it is very plausible that one can autonomously and wisely decide to defer to other people, and so one's autonomy is not violated. Another view advocates that, given our non-ideal world where social injustice exists and influences our epistemic judgements, we should not defer about moral matters because that harms those who lack privilege and perpetuates unjust social hierarchies (Dular, 2017; but compare with Wiland (2017) who argues that precisely because our society is unjust in certain ways there may be *pro tanto* reasons to defer to members of marginalised groups). Although this is definitely a matter worth investigating, and I am sympathetic to the point, I think there still is a question about whether moral deference would be problematic if our society were not unjust. I take this question to be worth investigating as well and I choose to focus on it because it will give us a broader insight into the nature of moral deference.

*based approach, the acquaintance view, and the authenticity argument.*⁵⁷ I will focus most of my attention on the first one, the moral understanding strategy, because it is the most developed view on the market. But the other arguments are also worth taking into consideration, so I will subsequently give a more concise assessment of those positions as well.

3. The Foes of Moral Deference

3.1. The Moral Understanding Strategy

Although the central element of this position is moral *understanding*, I include it in the *moral* strand of moral deference pessimism.⁵⁸ This is because what does the work in these arguments is a moral rather than an epistemic aspect. Hills (2009, p. 98, my italics) explains this: “Ethical traditions which defend trusting moral testimony and deferring to moral experts are, I will argue, missing something of vital *moral* importance.” The key point of this strategy concerns why moral understanding is important; and it is said to be most important because of its contribution to moral worth and virtue, among other things. This is the reason why I take the moral understanding strategy to be a moral instead of an epistemic variant of pessimism. As Lillehammer (2014, p. 113) puts it, this argument has at least one moral premise, so it is a moral argument to that degree, and it should “therefore be evaluated partly on [its] moral merits”.

The moral understanding strategy, as I will discuss it here, is a set of arguments that all focus on the value of *moral understanding*.⁵⁹ Its supporters admit both that moral knowledge can be transmitted through testimony and that one can enter

⁵⁷ Some of these strategies have been extended, more or less explicitly, beyond moral deference, to cover other types of deference, particularly concerning aesthetic matters. I discuss those proposals in chapter V.

⁵⁸ Not everyone does so, as they take it to be an *epistemic* version of pessimism, e.g. Howell (2014); Resiner and Van Weelden (2015). This does not matter, however, for the content of the discussion to follow.

⁵⁹ I take the supporters of this strategy to be Hills (2009; 2010; 2012; 2013), Callahan (2018), Hopkins (2007), and Nickel (2001). Hills has the most developed account, with Callahan bringing some revisions to it. Hopkins and Nickel do not discuss understanding as such, but it is clear that this is the notion they have in mind (although it is not quite as demanding as Hills’). They (Hopkins, 2007, p. 614; Nickel, 2001, p. 257) actually talk about it in terms of a *requirement* of morality, for belief respectively for action. Since Hills’ position is the most developed one I will take it as representative.

a trust relationship with a moral expert. Yet the suspicion surrounding moral deference does not disappear, and so they have argued that there must be some further norm that renders deference undesirable. There must be some other reason that explains our hostility to moral deference, and that is that it fails to give the agent moral understanding. What makes moral understanding so significant is that it is necessary for some things that are very important to us, as moral agents. We need it, first of all, in order to reliably do the right thing, as well as to be able to justify ourselves to others (Hills, 2009, pp. 106-107; Nickel, 2001, pp. 260-262). But, more importantly, we also need it to perform morally worthy actions and to be virtuous (Hills, 2009, pp. 108-119; Nickel, 2001, p. 257). I will focus here on the moral worth claim because it seems to me to be the most important and plausible one (and also because I discuss virtue at length in the next section). So, I will take it that the crux of the argument resides in this putative relationship between moral understanding and moral worth,⁶⁰ something that the strategists have not specifically addressed in detail. More precisely, we need to investigate how we can spell out claims like the following: “In order to act well, one *must* have a sense of the independent importance of various considerations, and act on the basis of that sense” (Nickel, 2001, p. 265, my italics), or:

Moral understanding is important not just because it is a means to acting rightly or reliably acting rightly, though it is. Nor is it important only because it is relevant to the evaluations of an agent’s character. It is *essential* to acting well.” (Hills, 2009, pp. 118-119, my italics)

By reconstructing the argument, we can see that there are two possibilities: a strong and a weak interpretation of this relationship. The former says that moral understanding is necessary for *any* degree of moral worth, while the latter that it is required only for *some* (presumably higher) degrees of moral

⁶⁰ At least for Hills (2009), this might apply more broadly than just in morality. She suggests, in further work (2017b) that there is an equivalent for aesthetic worth and aesthetic virtue (where aesthetic understanding plays the same role). Moreover, she argues for the value of understanding in general (2015), but it is not clear whether she would want to hold that there is an analogous understanding-based model of worth in any other domain. Hopkins has a structurally identical argument against aesthetic deference as well (2011), but it is also unclear whether he would want to go more broadly than that.

worth. In what follows I argue that the moral understanding strategists do not manage to establish that moral deference is problematic because they fail both to properly explain these key notions and to convincingly argue for the connection between them. I hold that the strategy ends up in a dilemma as any reading causes problems: either the premises do not stand (on the strong reading), or the conclusion is not significant enough (on the weak reading). I will start by briefly discussing the concepts of moral understanding and moral worth, and then I will reconstruct the moral understanding argument, ending with my criticism.⁶¹

As mentioned, the two key notions here are moral understanding and moral worth. Let us start with the former. Moral understanding is taken to be the grasping of the moral reasons that ground a certain moral claim, and the connection between these reasons and the respective claim. This is called *understanding why*, and is to be distinguished from both *understanding that*, i.e. comprehending a proposition and what it means (e.g. I understand that I ought to tell her the truth), and from *knowing why*, i.e. knowing a proposition and knowing the reasons that make it true (e.g. I know why I should tell her the truth: it is because I need to treat her as the autonomous and rational being that she is). According to Hills (2015, p. 663), what understanding-why has in addition to understanding-that and knowing-why is that it comes with (more precisely: it is partly constituted by) a certain intellectual know-how, a set of abilities which she names *cognitive control*. This means that the agent who has moral understanding will then be able to do things such as to follow an explanation of why *p*, be capable of giving such an explanation oneself, and to draw the conclusion that *p* and other conclusions in similar situations from the available evidence (Nickel, 2001, pp. 258-259 also discusses a very similar set of abilities).⁶² Moral understanding is taken to be the cognitive component of our responsiveness to

⁶¹ There are numerous objections that have been brought to the moral understanding strategy in the literature, see Callahan (2018); Groll and Decker (2014); Howell (2014); McShane (2018); Mogensen (2015); Sliwa (2012). There is no point in reiterating them so I will just add my own.

⁶² Hills (2015, pp. 666-667) explains that the extent to which someone has cognitive control is normally proportional to the degree of understanding-why that the agent has. Moreover, to show that the linguistic element (explaining why *p* in one's own words) is not always required, she also mentions that some of these abilities are specific to *explicit* understanding-why, i.e. do not apply to implicit or tacit understanding-why. The more intellectualised explicit understanding-why is her focus.

moral reasons (*orientation*, in Hills' terminology), which we use when forming moral beliefs and acting (Hills, 2009, p. 112).

Moral understanding, the philosophers contend, is important for many reasons. An agent who has moral understanding and not mere moral knowledge will be able to reliably do the right thing, will be able to justify themselves to other people, and, more significantly, will be able to act in a morally worthy way and achieve virtue. Out of all of these, particular attention is given to the connection between moral understanding and moral worth. The moral understanding strategists seem to hold that moral understanding is *necessary* for moral worth. But what is moral worth? The philosophers give a general and brief account of moral worth, so here I will be doing some reconstruction work, together with a very short description of how moral worth is discussed in general, outside of the moral deference debate.

First of all, moral worth, as a general notion, can apply, in some sense, to both to agents and actions, but I will focus on the latter here, to follow the literature. The moral worth of an action has been defined as “the extent to which the agent deserves moral praise or blame for performing the actions, the extent to which the actions speaks well of the agent” (Arpaly, 2003, p. 67). I will use, after Hills (2009), ‘acting well’ and ‘morally worthy actions’ interchangeably from now on. So what exactly counts as a morally worth action? One way to explain it is to say that for an action to be morally worthy it has to be performed *for* the right reasons (Hills, 2009, p. 113). There are two main views as to what constitutes the right reasons. Namely, the right reasons can be taken to be duty-based reasons or *de re* reasons. The first view says that actions have moral worth if they are performed out of the motive of duty (Kant, 1998; Stratton-Lake, 2000) and the second view wants the agent to act because they are motivated by the right-making features of the action (Arpaly, 2003; Markovits, 2010). For my purposes, no need to choose a side.

What all of this quite intuitively means is that in order to act well, we need more than just a right action; the action has to be done for that which makes it right. We would not say that a right action performed by someone by accident or, even worse, with the mistaken belief that it will bring about suffering, is a morally worthy action. Moral worth is the property of an action which helps us have more accurate evaluations of others as it eliminates (or at

least aims to eliminate) instances of doing the right thing as a result of extrinsic or resultant luck (i.e. doing the right thing by accident, for irrelevant or wrong reasons) and, on some views, even environmental luck (i.e. when one does the right thing, but could easily have not, had the situation been different).⁶³

Now what is more important for this discussion is the connection between moral understanding and moral worth. The problem is that, although we have a good grasp of what moral worth is, it is not so clear what role moral understanding plays when it comes to it. As mentioned, the moral understanding strategists use words like ‘essential’, ‘crucial’, and ‘must’ to refer to this role. So, plausibly, we can interpret the main claim of the general strategy as being that moral understanding is (at least) *necessary* for moral worth. That is, an agent will perform a morally worthy action only if they are doing so in response to its connection to the right reasons, appreciating them as grounding the action.⁶⁴ However, a further problem is that the strategists are also not explicit about whether one needs moral understanding to have moral worth *simpliciter* or only for some (higher) degrees. Accordingly, then, two possible interpretations of the argument, one strong and one weak, seem to be available. My aim is to see which, if any, reading is plausible. In what follows, I will first reconstruct the moral understanding argument, since the authors do not provide a very clear outline of it. My reconstruction will be intentionally ambiguous between the two readings precisely because the philosophers are ambiguous themselves. Subsequently, I will consider and critically evaluate the two interpretations, starting with the strong version and then going on to the weak one. In particular, I argue against P1 below.

Although the philosophers who support this strategy go about it in particular ways, fundamentally, the argument is this:

⁶³ A tangential point is that some philosophers think that at least some counterfactual strength is needed for attribution of moral worth. Arpaly (2003, p. 79) holds this, as she thinks that in addition to doing the right thing for the right reasons, an agent needs to have *moral concern* (which is equivalent to good will, responsiveness to moral reasons, a non-instrumental desire to do what is right) as well. A view that goes against the necessity of the counterfactual condition is Markovits’s (2010, pp. 210-214). However, this discussion can be put aside, as the focus here is on moral understanding.

⁶⁴ Hills (2009, p. 114) takes morally worthy actions to have a cognitive component (moral understanding) and a motivational component (good motives). So perhaps the most accurate characterisation of the view is that moral understanding is necessary for moral worth, but also necessary and sufficient for the cognitive component of moral worth.

P1: Moral understanding is necessary for morally worthy actions.

P2: Actions based on moral deference cannot be actions based on moral understanding.

C1: Actions based on moral deference are not morally worthy.

P3: Morally worthy actions are morally *pro tanto* better than mere right actions.⁶⁵

C2: Actions based on moral deference are morally *pro tanto* worse than actions based on moral understanding.⁶⁶

3.1.1. The Moral Understanding Strategy: The Strong Reading

I call the first version of the argument the *strong* reading because here the claim that moral understanding is necessary for moral worth means that it is necessary for moral worth *simpliciter*. That is, for an action to have *any* degree of moral worth, the agent needs to have, and to have acted out of, moral understanding. This is a plausible reading of the moral understanding strategy because from the way they argue, the philosophers seem to be saying that only knowing why a certain moral judgement is correct is neither necessary nor sufficient to act well. Instead, they emphasize the importance of grasping the relevant reasons, seeing the connection between the normative justification and the ought-claim, as well as of having the reasons guide one's judgements and actions. Moreover, they also discuss the importance of the abilities — the cognitive control — that can be gained only by having understanding-why, but not by having knowledge. All this, I take it, seems to say that merely knowing

⁶⁵ Or: mere right actions are *pro tanto* morally worse than morally worthy actions.

⁶⁶ I take this to be the best reconstruction of the moral understanding strategy. However, this is based on a 'worse than non-deference' claim about moral deference, which does not imply that moral deference is bad. If the supporters of this strategy want to hold the badness claim, then they would need to add further premises to their argument. For example, there could be a P4 of the form: 'It is always morally *pro tanto* bad not to perform morally worthy actions'. This would lead to the final conclusion C3: 'It is always morally *pro tanto* bad to perform actions based on moral deference (insofar as they are not morally worthy actions)'. An argument would be needed for P4, which seems at least *prima facie* vulnerable to worries about being too strong. But since I focus on P1 here, I will not settle the plausibility of this possible further premise. The strategists, of course, need not go for the badness claim, but stick with the 'worse than' claim. I do think the former is better, for reasons I have noted in chapter I, section 3, p 44, fn.. 45.

why a moral judgement is correct is not enough for moral worth. We *need* moral understanding.

Take, for example, Hills' (2009, p. 123; p. 117) claims: "If Eleanor fails to use her moral understanding, she is not properly oriented and she cannot act well"; "more [than good motivations and moral knowledge] is required for morally worthy action: you need to act for the reasons that make your action right"; "morally worthy action must be based on the agent's grasp of the reasons why the action is right and thus cannot be based on pure moral testimony" (Hills, 2013, p. 555). See also what Nickel (2001, p. 265) holds: "In order to act well, one must have a sense of the independent importance of various considerations, and act on the basis of that sense". Because of the textual evidence, P1 is to be read as saying that moral understanding is necessary for morally worthy actions of *any* degree or *simpliciter*.

Moral deference seems to take away the opportunity for morally worthy actions because it is not a means of transmitting moral understanding. What is transmitted through deference is moral knowledge and perhaps indirectly, at best, some understanding-that (if the judgement is accompanied by an explanation); but no understanding-why. When one forms a judgement or performs an action on the basis of moral deference it means that the reason for the new belief or action is the authority and expertise of the person to whom they have deferred. They believe or do something because a particular someone has said it is right, not because of the reasons that actually make it right. Moreover, because moral understanding is taken to involve a kind of knowledge-how, in the sense that it is typically related to the set of practical abilities aforementioned, it is something that cannot be passed on through testimony (Hills, 2009, p. 121). Just like one does not learn how to ride a bike or perform a surgery simply by reading or hearing other people talk about it, one is not able to navigate one's moral life by relying on deference. Thus, all of this gives us P2.

If moral understanding-why is necessary for actions to have any degree of moral worth, and moral deference essentially replaces moral understanding, then C1 — actions based on moral deference are not morally worthy — seems true. And if P3, i.e. morally worthy actions are morally *pro tanto* better than mere right actions, then we get to the main conclusion, namely that actions based on

moral deference are always *pro tanto* worse than actions based on moral understanding.

I take no issue with P2 and P3. P2 is true in virtue of how we use the notion of moral deference. That is, a belief or an action that is based (formed and sustained) on deference cannot, at the same time, be based on understanding. The sole explanation of the new belief or action is the expert's testimony and not the reasons that make it the right one. P3 is also at least *prima facie* intuitive, and most of the literature takes it for granted, so I will too. We evaluate each other differently in the light of our intentions, motivations, emotional responsiveness, and not just based on whether we have performed the right action or not. We appreciate and take as more valuable those actions which are done for the right reasons, which stem from an honest moral concern, rather than just being happy with right actions performed out of indifference, by accident or for the wrong reasons. Kant's famous example of the honest and the interested shopkeepers illustrates this in a very intuitive way. So I am happy to go along with P3. P1 is doing all the heavy lifting in this argument anyway.

I think that this strong interpretation of the general argument is the most consistent with the explicit positions of the philosophers who support this strategy.⁶⁷ However, I will now argue that on this strong reading P1 is implausible, and thus we have no good reason to accept the argument.

The first problem with P1 is that it is questionable whether an agent truly needs to have moral understanding in order to perform a morally worthy action of *any* degree. Recall that moral worth has been roughly defined as doing the right thing for the right reasons, and actions as having moral worth as long as the reasons that motivate the agent are the same as those which make the action right. Also recall that moral understanding is about a particular type of understanding, namely understanding-why. But why should we think that we need to understand why rather than to know why⁶⁸ an action is right due to some reasons? Why would it not be enough for at least *some* degree of moral worth to perform an action because of one's knowledge of the reasons that make it correct, even in the absence of an understanding of them?

⁶⁷ See also, for example, Hills' reply (2013, p. 557) to the idea that actions based on testimony could have moral worth: she says that the expert's action is morally worthy, but not the deferrer's.

⁶⁸ Contrary to Hills, Sliwa (2015) has argued that there is not such distinction at all: understanding can be reduced to knowledge.

According to Hills (2009, pp. 113-117) having just knowledge of why an action is right does not enable one to act *for* the relevant moral reasons. For example, when deferring to someone, one can find out what action one should perform, perhaps even obtain an explanation for that, and thus know why they should do that. Yet, the reason why one performs that action is the testimony of the expert, not the actual moral reasons that ground their testimony. Since moral worth just is doing the right thing for the right reasons, it means that such an action would not have moral worth:

Morally worthy actions are sensitive to moral considerations and for that you need to be oriented properly, not just in your outward actions but in your motivations, your choices, and your beliefs too. (Hills, 2009, p. 117)

This seems to follow if we are on board with the definitions of moral understanding and moral worth that the strategists use. But questions can be raised about why acting *for* the right reasons is limited to acting *directly* for the right reasons. Why not allow actions that are done indirectly for the right reasons, i.e. via deference, when the authority gives testimony that is based directly on the right reasons?⁶⁹ In other words, why should we believe these definitions of moral worth (rather than others, e.g. Markovits, 2010; Sliwa, 2016)? The response cannot be that when deferring you are not oriented properly, because orientation just is reasons-responsiveness and reasons-responsiveness encompasses moral understanding. But this is precisely the question: why do we need moral understanding for moral worth, when we could have access to the right reasons indirectly? Why have one over the other? We are owed a more extensive explanation for this.

Another response could be this. Hills (2015, pp. 670-671) argues that knowing-why does not give one the cognitive control that understanding-why does. More precisely, she emphasises this one ability that she takes to be specific to, and can be obtained only through, understanding-why, but not knowing-why: the ability to draw conclusions or give explanations about different but similar

⁶⁹ Sliwa (2016), for example, argues that moral knowledge is both necessary and sufficient for the cognitive component of moral worth, and moral understanding has no place there.

situations.⁷⁰ This ability translated into a kind of counterfactual condition for moral worth, at the level of cognition: would the agent still be able to figure out what the correct thing to do is, and why, in another but sufficiently akin case, if they did not have understanding-why in the first case, other things being equal?⁷¹ Hills seems to think that the answer is negative, and that is what gives understanding-why an advantage over knowing-why and, additionally, why we need the former for moral worth. I do not believe this is the case. Firstly, even if one does not understand why but merely knows why a certain action is the correct one, it does not necessarily mean that in a similar situation one would be clueless. Perhaps someone who understands-why would be quicker in judgement, or more reliable, but that does not mean that the one who only knows-why would certainly fail to handle a new case.

Secondly, assume that this is all true and that moral understanding is needed for moral worth and moral knowledge is not good enough. However, is it not implausible that this should mean *any* degree of moral worth? Moreover, it simply does not follow from the claim that moral understanding is necessary for moral worth that *no* degree of moral worth can be obtained when understanding-why is missing. Of course, we need not deny that moral understanding is a valuable thing. But to say that in its absence our actions have no moral worth seems too strong. Yet this is what the moral understanding strategists suggest:

If moral understanding is unavailable for you, there is no point in your trying to acquire and use it. Given your situation, *morally worthy action and proper orientation with regard to your moral beliefs are out of reach*. But if you are lucky and you have access to reliable people whom you can trust, you might be able get moral knowledge and, as a result, do the right thing. (Hills, 2009, pp. 123-124, my italics)

But compare the following actions: Aida gives a quarter of her income to charity because she believes that we ought to help others, at least when there is no significant cost to ourselves. Belinda gives a quarter of her income to

⁷⁰ Nickel makes a similar point (2001, p. 260; p. 264).

⁷¹ Whether or not such counterfactuals matter for ascriptions of moral worth is a controversial issue in itself (see, e.g., Arpaly, 2003; Isserow, 2019; Markovits, 2010).

charity because her friend tells her that she should, explaining that we ought to help others, at least when there is no significant cost to ourselves; Belinda wants to do the right thing, even if she does not understand why this is it. Cecilia gives a quarter of her income to charity because she wants to show off in front of her work colleagues. Doris gives a quarter of her income to charity by mistake, as she actually thought she was signing up for a raffle in which she could win an amount that is equivalent to a quarter of her income.

Out of the four women, only Aida acts for the right reasons, in the sense desired by the moral understanding strategy. She acts based on her moral understanding. Belinda does not have moral understanding, but she has knowledge of the right thing as well as knowledge-why of the reasons that make the action right. Cecilia does the right things for the wrong reasons, so she exhibits no understanding and no knowledge of the right reasons (perhaps she has them but does not use either). Finally, Doris acts rightly by accident: her action is not based on either moral knowledge or moral understanding of the right reasons.

According to the strong reading of the moral understanding strategy, only Aida's action has moral worth, while the actions of Belinda, Cecilia, and Doris are on a par, with no moral worth.⁷² But this is highly counter-intuitive. Belinda's action, even if perhaps not ideal, is different from the actions of Cecilia and Doris because at least she knows why what she is doing is right and is acting precisely on the basis of that, albeit indirectly, via deference. It should not be too controversial that her action is better, with respect to aspects that matter for moral worth (even perhaps just regarding the cognitive component), than the actions of Cecilia and Doris. Should we not then assign it at least *some* moral worth?⁷³ It should be widely accepted that an action like Belinda's is better than

⁷² This is particularly focused on the cognitive component of moral worth. Aida and Belinda are clearly different with respect to the non-cognitive motivational component: they *want* to do the right thing. But we can interpret Cecilia and Doris as not really having the desire to do the right thing. Still, regarding the cognitive component of moral worth, it seems that the moral understanding strategist is forced to say they are on a par. It is also implied that when you fail to have one of the two components of moral worth, your action simply has no moral worth.

⁷³ It is worth mentioning that some philosophers have argued that (at least some) actions performed on the basis of moral deference can actually be conceived as actions performed for the right reasons; in other words, that in some cases basing one's belief or action on the testimony of a good and reliable moral authority constitutes a right-making reason, e.g. Markovits (2010, pp. 218-219). Such arguments also constitute fair rebuttals of P1.

doing the right thing accidentally or for the wrong reasons, so it seems unreasonable to withhold a positive evaluation and not recognize *any* kind of moral worth in it (even if lower than if moral understanding would be present). Thus, I see no reason to believe the strong reading of the claim that moral understanding is necessary for morally worthy actions. Rather, we have strong reasons to deny that this refers to *any* degree of moral worth. Thus, the strategy in its strong version does not succeed.⁷⁴

3.1.2. The Moral Understanding Strategy: The Weak Reading

But there is another possible reading of the moral understanding strategy, a *weak* reading, which tells us that moral understanding is necessary for some, presumably higher, degrees of moral worth. This means that the strategy is to be understood as saying that one can have higher levels of moral worth only if one has moral understanding; or perhaps even that the more moral understanding one has, the more moral worth their actions have. In other words, that the degree of moral worth of one's action depends on the degree of moral understanding one has. Actions based on moral deference are not *highly* morally worthy, but they are not void of any degree of moral worth. So it would be *pro tanto* worse to defer insofar as that would put higher degrees of moral worth out of one's reach.

On this weak interpretation of the argument, P1 should be read roughly like this: acting out of moral understanding makes an action (at least all else being equal) more morally worthy to that extent. Moral understanding is not necessary for morally worthy actions *tout court*, but only in the sense that higher degrees of moral worth can only be achieved in the presence of moral understanding. This is a more reasonable interpretation because, as argued above, it is implausible to hold that moral understanding is necessary to have *any* degree of moral worth. Moreover, given that the strategists accept that moral understanding comes in

⁷⁴ Hills (2013, pp. 555-556) considers the moral worth worry from Markovits (2010). Her reply is very brief: she reaffirms her claim that only the testifier's action is morally worthy, while that of the deferrer is not.

degrees, they would probably also have to say that moral worth does too.⁷⁵ On this weak reading, P2 and P3 can stay the same. However, C1 must be modified accordingly and read as saying something more specific: actions based on moral deference are not highly morally worthy (or perhaps *fully* morally worthy actions). This intermediate conclusion would then only express the fact that moral deference precludes a certain level of moral worthiness, but it allows for lower degrees.

The moral understanding strategists neither explicitly hold nor deny this interpretation, though their formulations seem to be stronger. However, as argued, the strong interpretation seems implausible, so the weak version is the more reasonable and intuitive view. Even if some of the strategists would rather opt for the strong version discussed above, they would still be committed to this weaker one. This is because saying that moral understanding is necessary for moral worth *simpliciter* implies that moral understanding is needed for higher degrees of moral worth (unless one denies that moral worth comes into degrees). However, although this is the preferable interpretation, I hold that it is still unattractive because it gives us a conclusion that is not sufficiently significant, being not very informative. That is, it is not at all surprising.

But before I develop this worry, it is noteworthy that, if the argument is really to be interpreted in this way, the moral understanding strategist needs to answer some questions and provide further explanations. Where is the threshold for highly morally worthy actions to be set? Presumably, highly morally worthy actions will be actions performed on the basis of moral understanding. But how much moral understanding does one need for one's actions to be highly morally worthy? A stance also needs to be taken on when an action is morally worthy to being with. Is it when you act on the basis of knowledge-why or understanding-that (but do not have understanding-why)? Or do you not even need that, but only to indirectly act for the right reasons, via deference, for example? The strategist must offer a proper account of what kind of action counts as a morally worthy action, if this is the preferred interpretation of the view.

But assume they can do that. The weak reading of the argument tells us it is *pro tanto* worse to defer because deference precludes actions based on moral

⁷⁵ While the pessimists do not specifically address this issue, some philosophers from outside the debate have explicitly endorsed the idea that we can have degrees of moral worth, e.g. Arpaly (2003, p. 84); Markovits (2010, p. 238).

understanding, and having moral understanding is required for certain (higher) levels of moral worth. The problem with this view is that it does not tell us anything surprising. In this sense, the conclusion is not very informative. It is at least fairly obvious that one would have more moral worth if one would not defer, given that being able to do the right thing for the right reasons out of a global moral concern and orientation towards morality is quite plausibly morally better. It is also obvious that full, or high, moral worth is to be had by someone who can act in such a way. This is what our common sense conception of the virtuous exemplar is, after all; and we know that the closer one gets to it, the better; the more moral worth their actions have. So what the moral understanding strategy tells us is not that novel. In this sense, although the argument may tell us something true, it is not very forceful and not very interesting — as we already knew this to be true.

Moreover, it also overgeneralises. Recall the argument: moral deference is bad because it precludes us from obtaining moral understanding and moral understanding is necessary for high or full moral worth. If, by definition, moral deference is incompatible with high or full moral worth, then every single case of moral deference will be *pro tanto* problematic. But this seems implausible. Imagine a virtuous person who is properly oriented towards morality and usually does their own moral reasoning. But it so happens that in an isolated instance they need to defer. They have the correct moral orientation, the right motives, and good intentions. Is it reasonable to say that this is a problematic case of deference? It is true that the person's action is not highly or fully moral worthy, by definition. But they can reliably act in similar situations and are generally virtuous. At best, the strategist could say that this is *pro tanto* worse in the same sense that, e.g., the fact that we do not have even more good art in the world is *pro tanto* worse — i.e. in some sort of very general, and quite weak, sense. However, I would not think that this is a problematic case of deference. It seems too strong that a view would prescribe that every case of deference is a problematic case of deference.⁷⁶ So I hold that the weak reading of the moral understanding strategy is not the best pessimist view.

⁷⁶ There is also the question of whether our intuitions align with this view. Think of those people who are not only less than fully ideal but quite low on the virtue scale. It does not seem plausible to me to say that our negative intuitions about these people deferring relate to their actions not

To end this section, I also want to mention a general objection against the moral understanding strategy, understood in either its weaker or stronger form. This strategy does not manage to uphold the corollary claim of moral deference pessimism, namely the asymmetry thesis: there is a contrast between different types of deference such that some are more problematic than others. Even if the moral understanding argument would work, it would still not show that there is something particularly problematic about *moral* — or even normative — deference, something that we do not encounter in other subjects, since we obviously accept deference in the latter, but we are wary of the former. This is because the focus on the importance of understanding and its superiority over mere knowledge (Hills, 2009, p. 97; Nickel, 2001, p. 260) is not something specific to morality or other normative matters. Understanding is being recognized, more and more, as a significant epistemic value, something that might just be even more important than mere knowledge, in general (Grimm, 2012; Kvanvig, 2003; Pritchard, 2009). Hills (2017a) herself does extend her account of understanding-why beyond morality. If the asymmetry of deference is something the moral understanding strategists want to hold on to, then we also need an explanation of how they can do that given the above.⁷⁷

To sum up, as I see it, although it brings many insights into the moral deference debate, overall the moral understanding strategy seems to fail.⁷⁸ No matter how we interpret it, in a strong or a weak way, the argument ends up in a dilemma: it either makes use of premises that are implausible or it delivers a

being highly or fully morally worthy. They have bigger moral fish to fry, so to speak, such that the first thing that comes to mind would plausibly not be that their deference is bad because they failed to reach the highest moral peak.

⁷⁷ Perhaps they would do that by appealing to the notion of worth and virtue, which may apply only in some areas. But then there is the question of how widely that applies and whether it can adequately capture the asymmetry thesis as well. More on this in chapter V.

⁷⁸ Callahan (2018) builds upon Hills' argument by enriching the conception of understanding used, adding emotional and motivational components to it. Conceived like this, moral understanding would be sufficient (though not necessary) for moral worth and would put the agent on the path to virtue. But moral deference would be in tension with acquiring virtue, as deference discourages agents from seeking understanding. This is because someone with a settled view is less likely to try and gain understanding — but moral deference does exactly this, i.e. gives someone a settled view. I do not find this version of the moral understanding all that convincing either. First, there is no reason to think that moral *knowledge* is not sufficient for actions of at least some moral worth (even if it is not accompanied by the ideal emotional responses). Secondly, if the value of moral understanding comes from it being partly constitutive of virtue then at least some of my criticism of the virtue-based approach will apply. Finally, the idea that deference dis-incentivises understanding is not very strong because it will apply mainly (or only), as Callahan notes (2018, pp. 455-456), to those who lack the cognitive and temporal resources to further seek understanding. But that will then apply to a limited set of cases, and it is also not so obvious that deference would be even *pro tanto* bad in such situations.

conclusion that is not surprising. Moreover, it does not succeed in explaining the other claim that the strategists explicitly hold, namely the asymmetry thesis. The moral understanding strategy fails to build a good case for moral deference pessimism.

3.2. The Virtue-Based Approach

There are other views worth exploring. Although not as developed as the moral understanding strategy, a different version of moral deference pessimism moves from the moral worth of actions to the moral worth of *agents*, by focusing on virtue and character. According to what I call *the virtue-based approach* (Howell, 2014), there is always something sub-optimal about the agent who defers about moral (and plausibly at least some normative) matters because beliefs formed this way are isolated from their moral character.⁷⁹ According to this view, it is not the act of deference that is sub-optimal, but the agent and, more specifically, the agent's moral character. Deference exposes some existing flaws in the agent, but it also contributes to both immediate and long-term bad consequences for their moral character.

On this view (Howell, 2014, pp. 402-403), virtues are taken to be reliable dispositions that are characteristic of a person to act and feel in particular ways. The virtue of, say, generosity, will be manifested through generous actions as well as feelings of generosity, and joy in sharing and giving. A generous person will have reliable beliefs and intuitions about what situations call for generosity and how that should be enacted. All these aspects are unified and sufficiently integrated subjectively in the virtuous agent such that they reinforce each other. For example, virtuous feelings will give rise to virtuous actions and beliefs, and vice versa. By contrast, an agent can have the relevant dispositions to believe or act in certain ways — for example, in accordance to generosity —, but they may

⁷⁹ Crisp (2014) supports what he calls a phronetic version of moral deference pessimism, which seems to be a combination between the moral understanding strategy and the virtue-based approach. He holds that the virtuous agent, who has practical wisdom, will do the right thing for the right reasons, out of moral knowledge or moral understanding; if they do not, but they defer, they are morally and epistemically defective, and thus it is always *pro tanto* worse to defer about moral matters. His view is quite under-developed (being tentatively put forward in a response to Lillehammer, 2014), but what I say in response to the moral understanding argument and the virtue-based approach broadly applies to Crisp's version of pessimism.

not be integrated together. These dispositions would, then, not stem from virtue. An agent can perform generous actions without having the virtue of generosity. Virtue requires a subjective integration, understood as coherence, between beliefs, actions, intuitions, motivations, intentions, emotions, and various other cognitive and non-cognitive attitudes.

This view of virtues and the virtuous agent can then explain what is sub-optimal about moral deference in terms of the shortcomings that the agent has pre-deference, of those acquired immediately after deference, as well as those that will be long-term:

- Deference can indicate of a lack of certain virtues, not only because the agent does not know the target fact beforehand, but because the agent lacks either the ability or the drive to get the belief by non-deferential means.
- Deference, and the actions resulting from it, bypass moral character in that the beliefs and actions do not stem from the agent's present virtues (if they are present).
- Deference can result in an agent's having a virtuous belief, and performing a virtuous action, without the agent possessing the relevant virtue.
- Beliefs attained by deference can fail to integrate with the rest of the agent's present beliefs, and can fail to provide the proper ground for new beliefs. This prevents the agent from achieving higher degrees of virtue.⁸⁰
- Deference can either undermine, or frustrate the development of virtues insofar as the agent is likely to be in a poor position to feel or act reliably. (Howell, 2014, p. 403)

Based on Howell's formulation, I take point 1 to be about the pre-deference shortcomings; points 2 and 3 about the immediate shortcoming deference causes; and points 4 and 5 to be about the long-term effects of deference. The idea is that, first of all, the fact that an agent needs to defer to begin with shows that they have a flawed moral character and they lack virtue. However, deference leaves the agent in an even worse state. Because moral beliefs based on deference are isolated and disconnected from the other elements of the agent's character

⁸⁰ The charge of being uninformative that I have brought to the moral understanding strategy applies to this point too. It is fairly obvious (if not even conceptually true) that higher levels of virtues require people to be their own moral compass rather than to defer to others.

they will also preclude the agent from achieving the relevant emotional responses and other non-cognitive attitudes, the right motivations, and the appropriate intuitions. So moral deference will not allow the agent to strengthen their virtues or gain new ones, which in turn will make them less reliable and less able to apply their knowledge to new situations. This disintegration and lack of subjective coherence will also preclude the agent from achieving *full* or *complete* virtue, where we have the right beliefs, understanding, intuitions, attitudes, emotions, and motivations, with a reflective element. But, Howell (2014, p. 404) notes, not all instances of deference “will involve all or even most of these shortcomings. But together, these can account for much of our sense of what is wrong with moral deference”.⁸¹

Some sharp arguments have been brought against the virtue-based approach (Callahan, 2018; Hills, 2013; Mogensen, 2015). I will not reiterate them here, but instead I will add my own complaints. First, I think the focus should be on, as Howell (2014, p. 412) puts it, the crippling effects of deference (the long-term ones), rather than the pre-deference flaws deference indicates and the shortcomings to which it immediately gives rise. This is because they seem irrelevant in this context. The fact that sometimes moral deference triggers negative intuitions because it reveals there is something wrong with the agent pre-deference is, of course, an important point. It shows us that our reactions to moral deference might have various sources and that may change the way we think about its bad reputation. But does this matter when trying to explain what is wrong with moral deference *per se*? It does not tell us anything about moral deference in itself, or what moral deference does to the agent. It only tells us something about the pre-deference agent. Just like me baking a terrible cake does not tell us anything about baking. It only tells us something about me, namely that I am a terrible baker, and was a terrible baker before I baked this cake. But why would this matter if we want to evaluate baking? It seems to me that it does not.

The same goes for the immediate shortcomings that deference causes. Namely, according to Howell (2014, p. 403), the actions and beliefs based on

⁸¹ I am not sure whether this is meant to suggest that not all cases of moral deference are problematic, or that individual cases are in general not problematic and that we should only be looking at recurrent deference. Howell uses examples of both individual and recurrent deference and seems to take them all to be sub-optimal.

deference are not borne out of the agent's virtues, and these actions and beliefs can be virtuous without the agent actually having those virtues. However, this is, presumably, just a fact about what happens. They are not *effects* of deference, they are not shortcomings that deference causes immediately. So, again, this does not tell us much about moral deference and how it affects the agent. Moreover, unless we are thinking about the *perfectly* virtuous agent, it is perfectly coherent and plausible that, sometimes, even the (regular) virtuous agent will need some help. For example, someone might possess the virtue of generosity, but still encounter situations where it is really difficult to figure out what generosity requires of them and thus need help from someone who might know better. It is not inconsistent for virtue and deference to co-exist.

This is why I think focusing on the long-term consequences of deference is better suited for evaluating deference (*per se*). According to the virtue-based approach, beliefs formed on the basis of deference are not connected to our other existing beliefs and they do not give rise to the virtuous dispositions to act and feel in appropriate ways. This will make the agent's moral reasoning more difficult, as it will be harder for them to achieve reflective equilibrium and a coherent moral outlook. Because of this, virtue is not reinforced and developed, and the agent is also precluded from achieving higher levels of virtue, of the reflective kind.

However, this view of moral psychology is a bit odd. It does not seem to be based on regular people actually deferring, but on exaggerated examples; this is precisely what leads our arguments astray, as I have argued in chapter I and in this chapter. Are beliefs and actions based on deference necessarily isolated from our character?⁸² They seem to be if you think of characters like Gary, who never trusts himself to make a moral decision so he always defers to Google Morals (Howell, 2014, p. 401); or Urkel, who has a good understanding of ethics, but has no moral intuitions whatsoever. He always uses Google Morals to get the answer to his moral questions, although after that he can come up with the right explanations (Howell, 2014, p. 398).

⁸² Hills (2013, p. 557) very briefly considers such an objection too: "although a moral judgement made on the basis of testimony could well be cognitively isolated, first, it certainly need not be. That helping the needy is morally right could fit right in with all my other moral judgements." I agree with this, but my point will go even further.

But such agents are not representative. For one, their pre-existing flaws are so deep that they already guarantee that they are not even able to have the subjective integration of character that Howell wants and purports to be the defect of deference. Of course the person who is unable to have any moral intuitions will not suddenly acquire intuitions when deferring. And of course the person who does not trust himself at all will not suddenly acquire virtuous motivations, intuitions, and feelings when deferring. By stipulation they will have those defects. So, it seems to me that using such cases to show that deference is sub-optimal is not fair. All the characters that Howell describes are so flawed, from the beginning, that they cloud our intuitions about deference; we simply cannot tell. Second, using these examples disregards what I take to be realistic cases of deference. That is, cases of people who are concerned with their moral questions, who perhaps are worried about the seriousness of the consequences of their actions; people who care about doing the right things and people who are flawed just as much as we all are. If we think that this type of cases is the right one to focus on, then we can see that the original claim is not so plausible in this context. Even if such agents end up deferring, they would plausibly try to integrate their new beliefs or actions as much as possible simply because they care, to a smaller or larger extent, about morality. It is implausible to think they would just let this new belief float around unattached to anything else; it will not be a shocking new thing that is unconnected to other beliefs. It will certainly not be perfect: they might not acquire intuitions, feelings, or understanding-why (although it is not impossible they would). But this does not amount to isolation. The new belief will connect to other beliefs, in search of coherence; and, if the purpose was to find out what to do in a certain situation, it might well give rise to some level of motivation to perform the action in question — after all, this is why the person wanted to defer in the first place. This picture of deference leading to cognitive isolation makes sense only if we imagine agents who have no clue about the topic they are deferring to begin with, who do not get any explanation for the judgement deferred on, and who do not make any effort to consider the new belief or action in the bigger picture of their moral outlook.

Now, to be fair, Howell (2014, p. 402; p. 404; p. 405) says at various points that deference does not *guarantee* integration. But does any way of forming beliefs *guarantee* this total integration of beliefs, feelings, intuitions, and

motivations? Even doing our moral reasoning ourselves can leave us unmotivated at times; I can also certainly come to a moral conclusion on my own and yet do it coldly such that I do not attain the appropriate moral sentiments. If we are being honest, it is very difficult for an imperfect moral agent to acquire such integration, even if left to their own devices.

You might want to say that, nevertheless, this lack of integration is bad. It does not, for example, allow us to achieve “a complete, more reflective virtue” (Howell, 2014, p. 406), which Howell does say about deference. However, if the view boils down to this, then it becomes uninformative because it is unsurprising. As I have argued in the previous section, it is well-known that to be perfectly virtuous, you have to have the correct moral beliefs, feelings, motivations, intuitions, understanding, and so on. This is precisely how the virtuous exemplar is supposed to be, conceptually. But this is not a very interesting view of deference then, because we already know this. Saying that deference about moral matters is not ideal, with respect to complete virtue, is almost trivial.

Moreover, there are other questions that need to be answered about whether we should even aim for that. Remember Wolf’s (1982, p. 421) claim from above:

For the moral virtues, given that they are, by hypothesis, *all* present in the same individual, and to an extreme degree, are apt to crowd out the nonmoral virtues, as well as many of the interests and personal characteristics that we generally think contribute to a healthy, well- rounded, richly developed character.

So perhaps we might not want to be the kind of people that moral saints are, even if that might be, in some sense, good or even best. Full and complete virtue might not be worth pursuing after all, and it is unclear whether we would be justified in faulting people for not doing so.

3.3. The Acquaintance View

Lord (2018) also puts forth an argument for why deference is defective, which is meant to cover both the aesthetic and the moral domains. His view is

centred on the concept of *acquaintance*, which is required for appreciative knowledge, the kind of knowledge that makes appreciation possible. Acquaintance is defined “as a kind of direct access to an object, property, or fact” (Lord, 2018, p. 89). For example, seeing counts as acquaintance. But imagination and having certain affective states can also get us acquainted with some properties, according to Lord. Imagining how a piece of music goes or experiencing a feeling like repulsion in response to a situation, can also give us direct access to the normative features of an object (Lord, 2018, p. 91).

Being acquainted with the properties of an object⁸³ means that one is in a position to rationally have certain reactions to that object. For example, take Hanna who sees the bust of Nefertiti in the museum. The fact that she saw that the bust is beautiful gives her both knowledge of the fact that the bust is beautiful and it enables her — and makes it rational for her — to appreciate the bust. This is called *appreciative knowledge* and is the kind of knowledge “that allows one to fittingly have the full range of affective and conative reactions” (Lord, 2018, p. 77). Hanna possesses the fact that Nefertiti’s bust is beautiful as a reason to react in certain ways towards it, e.g. forming certain beliefs, desires, feelings, intentions. More precisely, “Hanna is put in a position to manifest knowledge about how to react affectively and conatively to the bust’s beauty” (Lord, 2018, p. 79). These reactions, Lord thinks, are to some very particular features of the bust, features that the bust’s beauty depends on. Given such aspects are so fine grained, it would be very difficult for us to be able to react in the appropriate way if we were not acquainted with the object ourselves and instead only heard a description of it.

For contrast, take Clara, who asks Hannah about her trip to the museum and defers about Nefertiti’s bust being beautiful. She comes to have the knowledge that Nefertiti’s bust is beautiful, but she cannot have the full range of reactions to all the reasons that the aesthetic qualities of the bust give us. At best, Clara can have a reason to go to the museum to see the bust herself. But it is not rational for her to be in awe with the bust, to feel admiration for it, or hope that it will be preserved forever. She does not have appreciative knowledge, as she does not possess the relevant facts as reasons for appreciation.

⁸³ In contrast with the literature before him, Lord focuses on being acquainted with the properties of an object, and their distributions, rather than the objects themselves (Lord, 2018, p. 89).

This is what is problematic with deference in Lord's view. Testimony is not sufficient to enable us to gain appreciative knowledge as it cannot make us acquainted with the specific normatively relevant properties. Being acquainted with moral facts allows one to possess them as reasons and then have appreciative knowledge, which is the basis for the fitting conative and affective reactions. Moral deference cannot make it possible for us to possess these facts as reasons, which limits what we are able to do with the normative information that we have. This applies to the aesthetic and the moral case, and perhaps more widely regarding other normative matters.⁸⁴ Lord thinks, however, that none of this tells us we *ought not* to defer. Deference can still give us knowledge and it can also permit us to possess reasons for *some* reactions, even if not for all. Nevertheless, according to Lord, appreciative knowledge is central to our moral and aesthetic lives and to being good moral and aesthetic agents; since deference does not give us appreciative knowledge, it is, in this sense, defective.

The main problem I see with this account of the badness of deference is that it is extensionally inadequate. More precisely, it is underinclusive. It does not account for two main sets of problematic cases of deference. The first includes instances of deference regarding how to weigh up the considerations are relevant for a given decision about what to do (or believe). In such situations, it seems that a person is acquainted with the relevant normative properties because they are literally in a moral conundrum themselves, reflecting on those properties in the process of trying to make a decision. It is just that they do not know which are more relevant or how to weigh them up, and so they defer. Generally, the intuition is that such cases of deference are also problematic. But, on this view, why would they be, if the person does not lack acquaintance? The view predicts that such a case of deference is not problematic.

But why should we think that one is acquainted with the normative properties in this situation? One might worry that Lord's notion of 'acquaintance' is much more substantive, more epistemically loaded. For example, he often says that acquaintance is the kind of things that puts us in a position to use certain facts as reasons for a range of reactions. However, he also defines it like this: "I understand acquaintance in a broad way to include any kind of direct access to

⁸⁴ There are some problems that might arise in the moral case (e.g. how exactly we get acquainted with moral facts), but Lord (2018, pp. 88-92) anticipates quite a few of them. His responses are highly controversial, but for the purposes of this discussion I will accept them.

an object, property, or fact” (Lord, 2018, p. 89). This can happen, he says, by encountering real moral situations or even by imagining such situations (Lord, 2018, pp. 92-93). But what I have in mind is exactly this: a real life situation that someone is in, where they are reflecting on what they should do.⁸⁵ The person would probably know how to use certain facts as reasons, but they do not know what to do all-things-considered and how these reasons weigh up against each other. If this does not count as acquaintance, then I do not really know what acquaintance is.⁸⁶ Denying this would be quite implausible (and would make the concept of acquaintance strangely stipulative). The person in my example is in a position to use certain facts as reasons but cannot weigh them up to see which one(s) prevail(s). So, we should not deny that they are acquainted with the normative properties, but only that they cannot manifest appreciative knowledge, for whatever reason (e.g. uncertainty is not that exotic). This suggests that acquaintance is not sufficient for appreciative knowledge, and so an explanation based on acquaintance is unfruitful, if what truly matters is appreciative knowledge.⁸⁷

The second set of problematic instances of deference that this view does not cover involves cases where someone defers about a piece of knowledge that can only be obtained through means that do not allow for acquaintance anyway. I have in mind here moral knowledge gained through moral theorising.⁸⁸ Lord (2018, p. 93) allows that certain ways of gaining knowledge will not give us acquaintance or appreciative knowledge:

Investigating necessary moral truths has always been part of ethical theorizing. Indeed, it has monopolized most theoretical discussions. The pursuit of contingent ethical truths is not popular among moral philosophers (even applied ethicists). Our theoretical priorities, then, suggest that in ethics the necessary truths take

⁸⁵ Not only is someone experiencing, first-hand, a concrete situation where they have to make a decision, but they are also reflecting on their case. About this Lord (2018, p. 94) says the following: “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. We become acquainted with these properties via these imaginative experiences.” If imagining a case gives us acquaintance, could we deny that living the case, so to speak, does not?

⁸⁶ Lord’s (2018, pp. 81-82) Akshai example seems to support this.

⁸⁷ But Lord may accept this, as I will explain next, when anticipating his possible reply to my next objection.

⁸⁸ The same should hold for moral inferential knowledge, if it can be taken as significantly different from moral theorising.

precedence. How we learn about them is a matter of great controversy. It is not popular, however, to think that we learn about them via acquaintance. If this is right, then it's easy to see how armchair knowledge can threaten my project.

His response to the challenge is the following:

The problem is that our theoretical priorities give a misleading picture of how most moral learning works. It is implausible to think that knowledge of necessary moral truths plays a prominent role in the average person's acquisition of new moral knowledge. I doubt this is true even of moral philosophers. I myself do not flat-out believe many necessary moral claims. A large part of the theoretical interest in such claims is that they are hard to learn! Thus, I don't think we should allow moral philosophers' theoretical priorities to shape how we think of moral epistemology. (Lord, 2018, p. 93)

Here Lord does not seem to deny that moral theorising does not give us acquaintance. Nor does he say that it is problematic or defective as a source of knowledge because of that. He only claims that this is not so important given that not a lot of our moral learning occurs this way — it is not the paradigmatic way of obtaining moral knowledge. However, it is not nonexistent. But then I would think that moral deference about some piece of knowledge obtained through moral theorising (however rare that is) should not be defective either.

Take an example that Lord (2018, p. 93) mentions: judgements about the modal status of some moral truths. Let us say that Jane did some moral theorising about the modal status of some moral truth. She came to the conclusion that it is *necessarily* true that murder is wrong. As Lord admits, this moral theorising does not come with, or does not happen through, acquaintance. So why would it be problematic for me to defer to Jane regarding the judgement 'it is necessarily true that murder is wrong'? What makes moral deference problematic is precisely that it does not get us acquainted to the relevant normative properties. But if there was no acquaintance to begin with, then what exactly is the issue? No matter how rare such cases are (in virtue of the purported rarity of armchair moral theorising about necessary moral truths), they can still occur. And if they do, this account does not have the resources to explain their fishiness.

However, Lord may have an answer to this objection. Acquaintance, at least in the moral case, may not be necessary for appreciative knowledge (2018, endnote 18):

Note that there are two crucial elements of my view. The first is the claim that deference lacks certain rational powers because it doesn't provide appreciative knowledge. The second element is the claim that deference doesn't provide appreciative knowledge because it doesn't acquaint us with the normative properties. One could—and I would—accept the first part of the view even if the second part is false. Thus, even if acquaintance isn't required for appreciative moral knowledge, it could still turn out that deference is defective because it doesn't provide appreciative moral knowledge.

So, the problem is actually that the deferrers in my examples do not exhibit appreciative knowledge. But now a question arises: why does deference not give us appreciative knowledge? What do the other ways of gaining knowledge have that deference is lacking, if the issue is not acquaintance? As it stands, Lord's argument seems to heavily rely on acquaintance. So we need a story here regarding how the view goes if we take acquaintance away. Moreover, if we do take acquaintance away from the explanation of moral deference, Lord's explanation would lose (at least some of) the unified character it purports to have in virtue of wanting to give the same argument for the problematic character of both moral and aesthetic deference (the latter being discussed in chapter V). This, of course, is not a knock-down argument against Lord's view because it is simply asking him to say more about his position. However, given that his position, as it stands, is based on acquaintance, it is reasonable to suspend judgement about its success until we see how it would look when acquaintance is taken out of the picture and replaced with a new argument.

3.4. The Authenticity Argument

The authenticity argument represents a variety of moral deference pessimism that has not been fully developed yet (Mogensen, 2015; Skarsaune, 2016). It is worth mentioning, however, as another option on the table. The

philosophers that propose this position focus, of course, on the value of authenticity. The explanation for the problematic character of moral deference is, basically, that it undermines one's authenticity, which is important for oneself and for other people with whom one interacts.

Authenticity is roughly defined as a “harmony between the inner and outer facets of a person” (Mogensen, 2015, p. 276). In other words, a person is authentic if their actions express their true beliefs, feelings, and motivations; their true self or who they really are. A person acts authentically if they are guided by their own understanding of the reasons that are relevant for that action (Skarsaune, 2016, p. 352).

One argument for pessimism about moral deference is that it simply follows from the demands that authenticity places on us. Authenticity places demands on us because, at least in our culture, it is a significant value to which we have a particular attachment.⁸⁹ This attachment is precisely what makes it valuable and what makes it such that we ought to figure things out for ourselves.⁹⁰ Authenticity demands from us that we manifest our true selves and that we achieve that harmony between our inner life and behaviour. The attachment to this value is specifically important in morality because our moral beliefs are central to our identity, as “we are defined by our values and ideals” (Mogensen, 2015, p. 277).⁹¹ Deference prevents us from manifesting ourselves as we truly are. Forming beliefs and performing actions based on testimony are not expressive of one's own ideals, desires, emotions, or values — they express someone else's values, namely the expert's. As such, it makes us inauthentic.

Another argument against moral deference is focused on the importance of authentic behaviour in human relationships in general. We want people to behave authentically towards us so we can assume they want that from us too.

⁸⁹ The important question of why people value authenticity, and whether they should, remains an open question (Mogensen, 2015, p. 276).

⁹⁰ Mogensen (2015, p. 275) flips the definition of pessimism, saying that it is mainly a view about the importance of figuring things out for ourselves. This way of framing pessimism really hits the spot. Virtually all pessimist views fundamentally argue against deference by showing what we are losing when we defer; i.e. by arguing for what the value of figuring things out on our own means. After all, this is what causes so much trouble: we do not think that deference *qua* deference is problematic, but that deference is problematic in virtue of obstructing, in one way or another, the value of figuring things out on our own. Thus, pessimism can be seen, in general, as following from why each of us thinks figuring things out on our own is important.

⁹¹ Mogensen (2015, p. 278) relies on some limited empirical studies for this claim. The question is also open regarding the conclusion of these studies and whether moral beliefs are indeed central to our identity.

What we all want is that we engage in meaningful, authentic interactions. Imagine, the story goes, a city where all the people decide how to act in every human interaction by deferring to a highly reliable Google Morals. These people do not understand why they act the way they do and they definitely do not do it out of a recognition of the right reasons; they simply defer to Google Morals all the time and learn, in that way, to do the right thing. But it seems that we would not be happy with such people. Even if the world would be safer than ever, even with the guarantee to be always treated fairly and nicely, the intuition is said to be that we would be missing something. The thing that is missing is precisely the value of authentic interaction, about which we seem to care a lot.

One explanation for why authentic behaviour is important is that it shows something about the people we come into contact with: their normative judgements, the fact that they have moral understanding, and how they rank their values. We generally want to know this because we want to know who these people really are. The second reason is that it provides a genuine personal engagement with others; in the absence of an authentic behaviour, “there is a sense in which you are not fully interacting with her as an individual at all” (Skarsaune, 2016, p. 353). This is why the more personal and the more high stakes a situation is, the more valuable authenticity seems. Moral deference takes this away from us because it puts us in a position where we cannot authentically interact with one another.

Skarsaune’s argument is a direct response to Enoch (2014), whom I will discuss in the following chapter. Here it suffices to say that Enoch (2014) presents an all-things-considered optimist argument. He argues that in certain cases, where there is a risk of doing wrong to others, we ought to defer if there is someone who is more likely to get it right than ourselves, all else being equal. There is nothing that can outweigh minimising the risk of wrongdoing such that non-deference would be preferable. Skarsaune responds by saying that the value of authentic interaction is actually such a value which can, at least in some very specific cases, outweigh it. He argues (2016, pp. 355-356) that we would all-things-considered want others to behave authentically rather than inauthentically, even if that means that we risk being wronged in some way, if two conditions are in place. One, if we are engaging with people with whom we are in close personal relationships; and, two, if the stakes are low or moderate.

The basic point seems to be that, based on our intuitions, sometimes we care more about authentic interaction and other times we care more about minimising the risk of wrongdoing. Depending on whether the two conditions from above are met, sometimes the value of authentic interaction will outweigh minimising the risk of wrongdoing and other times it will be the other way around.

The first general problem with the authenticity view is that it assumes that deference is incompatible with being authentic, and thus it does not manage to explain the fishiness of certain cases of deference. But authenticity need not necessarily clash with moral deference. Anyone who defers can do so in an authentic manner, namely in a way in which they express their true self and what they hold dear; for example, the desire to do the right thing and the motivation to be a good person. Asking someone for moral testimony can represent an explicit recognition that we are limited and we know others know better. We do not need to have all the right answers to be authentic. Moreover, deference need not come hand in hand with a pretence that our newly acquired moral beliefs have originated in us. We can always be honest and admit all of this to the people with whom we interact. We can also even defer to someone who has the same values as us, or who uses our values to help us come to a decision (Howell, 2014, p. 391-392; McGrath, 2009, p. 322). Thus, our behaviour would be representative (albeit indirectly) of our actual inner world view. If all of this is true, then deference in such cases does not undermine authenticity and so, on this view, it should not be problematic. This makes the view underinclusive because deference does seem problematic even in cases where we are authentic and honest about our situation, intentions, and motivations.⁹²

Another problem is that the view might not do well tested against cases of bad people. Take, for example, sexist or racist people, who can answer their own moral questions (albeit incorrectly when it comes to issues related to sexism and racism). Would we want them to be authentic and express their true racist and sexist selves or to defer and at least do the right thing? Of course, having them pretend to be something they are not is not great, for many reasons, but it

⁹² Mogensen (2015, pp. 263-264) is one of the pessimists who explicitly concede that many cases of deference are acceptable (but it is not clear if he means that there still is a sense in which they are *pro tanto* worse). He seems to restrict the range of cases he is interested in to those where the person could have figured things out on their own. But even in those cases what I say above can apply.

is intuitive that we should prefer a world where at least people know they should not be explicitly racist and sexist (even though they secretly might be) to a world where they think they can explicitly be like that. Intuitively, we would (at least all-things-considered) want such people to defer rather than being authentic and express their own world view. Can Mogensen's and Skarsaune's views give this verdict?

Mogensen would say that someone who is sexist, call him Robert, would be subject to the demands of authenticity and as such he has a duty to figure things out on his own. The duty applies to him because he is able to figure out the answers to his moral questions alone. However, if this duty is *pro tanto*, then Mogensen could say that it is outweighed by the value of Robert learning what the right thing is by deference or by the duty to minimise the risk of wronging others with his sexist behaviour. If this is the case, all is fine for now.

For Skarsaune it might get a little complicated. He claims that:

in an impersonal relation and a decision where the practical stakes are high, we prefer deference (when that improves the odds of right action), but in a personal relation and a low-stakes decision, we prefer authenticity. (Skarsaune, 2016, p. 355)

We can imagine a situation where Robert is my colleague. I do not know much about him, but my interactions with him have been nothing but pleasant. I even consider that we might become friends, as I do not know he is sexist. In a conversation at lunch, I tell him that in my home country some people often comment that women lose their right not to be harassed if they do not dress conservatively enough. I ask him what he thinks about that, sort of jokingly, as I expect him to see it as an absurdity, like I do. But just then he is called to his desk and he has to go. We say we will continue the conversation at lunch tomorrow. Later in the day, he meets his friend, Lisa, and tells her about our conversation. Lisa knows Robert grew up in a sexist family and environment, and although he tries to reform himself, he still has trouble identifying sexism. She tells him that of course it is wrong that women are harassed, no matter how they dress. Let us assume that Robert does not see how the claim is grounded, but nevertheless he defers to Lisa, and when we meet for lunch the next day he

tells me he believes that it is wrong that women are harassed, no matter how they dress.

The relation is personal enough, as I consider becoming friends with him, but the stakes are also low enough — it is just a lunch conversation. Finding out that he would agree with my co-nationals would not harm me that much⁹³ because, after all, I grew up with some such people, and sexism is, sadly, old news to me. But according to Skarsaune’s account, it would seem that Robert should not have deferred, but rather all-things-considered be authentic with me. This seems like the wrong result. Robert did well, all-things-considered, deferring. At least that way, he will learn what the correct thing to believe is. We probably want these people to defer precisely to start (or to continue) their journey to being better people. Moreover, maybe they themselves want to be better and to bury their authentic, sexist selves, and get help in becoming virtuous people. Getting rid of internalised sexism is, however, difficult and requires a longer process, which will also plausibly involve some deference. In that process, the person might not want to express their true beliefs and feelings to others, exactly because they want to change them. It does not seem like there is anything wrong with this. It is not clear why authentic interaction would be preferable in such a case.⁹⁴

But assume that Skarsaune would agree and say that, in such a case, the value of authentic interaction is indeed outweighed, and it is all-things-considered better to defer. So he would be in the same place as Mogensen, thinking that it is *pro tanto* bad that Robert deferred, but it was, overall, the appropriate thing to do. However, this still seems strange. In what sense is it *pro tanto* bad that Robert did not reveal himself to me as the sexist that he is? What value does an authentic interaction have when people are (at least partly) bad? There does not seem to be a sense in which I should wish that Robert was authentic with me. We could think that, if I knew his true self, I would not want to become friends with him anymore; in this sense, it would be good for me to

⁹³ I also do not think that finding this out would wrong me as such. What would wrong me (or harm me) is him acting in a sexist manner towards me.

⁹⁴ This is on top of the fact that Skarsaune (2016, pp. 355-356) does not have answers for mixed cases, where the relation is personal, but the stakes are high, or the relation is impersonal and the stakes are low, as he says intuitions differ when it comes to whether authenticity can outweigh other values in such situations. But this means that this view is unable to handle all cases of moral deference, and so it is underinclusive.

know his authentic self. That could be true. But if he is genuinely trying to change, and he has a hard time telling me about his past — after all, I am just a colleague — then him being authentic with me would not help any of us. It would just ruin a potential friendship. So it is hard to see the *pro tanto* value of authenticity when it comes to people who have (serious) character flaws.⁹⁵

One thing the supporters of this argument could do is to say that to hold sexist values is inauthentic, in some sense. This is not addressed, or even implied, in Mogensen and Skarsaune, but they could fashion such an appendix, on the model of views that take integrity as moral purpose, where:

objective integrity requires that agents have a sure grasp of their real moral obligations (Ashford 2000, 246). A person of integrity cannot, therefore, be morally mistaken. (Cox et al., 2018)

But this would make the authenticity argument very narrow with respect to the kinds of cases of moral deference it could explain. Most, if not all, of us, are at least sometimes morally mistaken and we do not have a full grasp of the real moral obligations. So that might make all of us inauthentic on this strong reading. A weaker reading of authenticity could be developed, but a threshold, that is neither too high nor too low needs to be set, and then all tested against different cases. This would require extensive development of the authenticity view. Until then, the view is subject to these problems.

4. Concluding remarks

The existing literature on moral deference offers a lot of insight into the problem of moral deference. But the accounts that have been put forth are not without problems, as I have argued above. So, I think we can aspire to do better.

⁹⁵ Assume, however, that Robert is not trying to change. He is just sexist. Would it be *pro tanto* bad if he did not reveal himself as such to me, but deferred? My intuition is there is something bad about him not telling me the truth, but I do not think it has to do with authenticity. Rather, I do not like the fact that I was lied to. Deference does not even matter in the discussion. He could have reasoned his way to lying alone. You know the familiar story: ‘I know that political correctness has taken over the world, so I cannot speak my mind anymore. She definitely expects me to say that it is wrong for women to be harassed. So, I will, because otherwise I could lose my job.’ But then I am not sure that the demands of authenticity are doing the work here. Rather, it seems like the demands of honesty matter. Are they the same thing? The authenticity argument might look different, if we would say that authenticity simply reduces to honesty.

In what follows I will show that even though the current pessimist views have not been successful, it does not mean that moral deference optimism has been established. In the next chapter I argue that the optimist arguments presented in the debate are actually compatible with pessimism, such that we can still pursue the line of thought that moral deference is problematic.

Chapter III: Friends of Moral Deference

In this chapter I discuss the arguments that moral deference optimists have put forward and argue that they are not successful in showing that pessimism is wrong. Since my own version of pessimism was not on the market when these views were developed, I will discuss the existing landscape first — i.e. existing views of optimism versus existing views of pessimism — and then address how my view fits in.

Moral deference optimism is, simply put, the denial of moral deference pessimism.⁹⁶ To refresh our memory, the best interpretation of pessimism, as far as I can tell, would be the following: for any case C in a certain range R (to be specified), moral deference is *pro tanto* bad in C, even though it may be all-things-considered better to defer in C (and in many other cases, too).⁹⁷ Denying moral deference pessimism sometimes takes the form of literally arguing against specific arguments that pessimists make (Groll and Decker, 2014; Lillehammer, 2014; McShane, 2018; Reisner and Van Weelden, 2015; Sliwa, 2012; Wiland, 2017). However, I will not be concerned with that here because I want to tackle the general optimist approach. I take it that there are two main strategies within the optimist framework. One is to argue that *moral* deference is not generally problematic in itself or *per se* (Sliwa, 2012; Groll and Decker, 2014). The other one is to argue that moral deference is sometimes good (Sliwa, 2012; Driver, 2015), better than non-moral deference (McShane, 2018; Wiland, 2017), or morally required (Enoch, 2014). I will call the former *the negative strategy* and the latter *the positive strategy*, and in what follows I will deal with them in turn.

My aim is to offer a general response to moral deference optimism, so I will not criticise the minute details of the views presented. I want to show that even if these optimist arguments work perfectly as they are, they do not successfully make the case against the best version of moral deference pessimism. This is, first, because both the negative arguments against pessimism and the positive arguments for optimism can be accommodated by the most

⁹⁶ There are philosophers who argue for epistemic optimism, namely for the claim that moral experts exist and that we can reliably identify them (Cholbi, 2007; Driver, 2013; Jones and Schroeter, 2012). I will not engage with these accounts here because, as I have said, this thesis is not concerned with the epistemology of deference. I thus focus on non-epistemic optimism.

⁹⁷ Where the range of cases R can be limited to cases where the deferrer could have made the judgement on their own or expanded to cases where they could not have done so.

plausible pessimist views. In other words, the respects in which moral deference is unproblematic according to these positive arguments are compatible with the respects in which pessimists (should) regard moral deference as problematic. The sorts of views that get advocated as forms of moral deference optimism are not incompatible with the best version of moral deference pessimism. Secondly, the existing optimist arguments do not show that *my* version of pessimism is false. Namely, pessimism under the claim that *recurrent* moral deference is always *pro tanto* bad. As such, optimism is unsuccessful in fully redeeming deference.

1. The Negative Strategy

According to this set of arguments (*debunking arguments*, as Hills, 2013, p. 558, calls them), which I have anticipated in the chapter I and II, there is no general problem with *moral* deference in itself or *per se* (Sliwa, 2012; Groll and Decker, 2014). This is not to deny that there are some instances of moral deference that are problematic. But those cases, the optimists say, “aren’t problematic because of any general problem with moral testimony” (Sliwa, 2012, p. 177). Rather, they can be explained by factors that are dissociable from moral deference *per se*. Typically, those factors also explain why some instances of non-moral deference are problematic, which shows, according to the optimists, that the asymmetry thesis — at least understood as “a principled difference between moral and non-moral testimony which renders the first, but not the second, problematic” (Groll and Decker, 2014, p. 54) — is not true.⁹⁸ So I take the general argument to be something like this:

⁹⁸ The asymmetry thesis is discussed in detail in chapter V.

P1: We have intuitions that suggest that moral deference is problematic *per se*.

P2: If we have intuitions that moral deference is problematic *per se*, then, other things being equal, those intuitions are best explained by there being something problematic about moral deference *per se*.

P3: All these intuitions can be explained away by factors that are not features of moral deference *per se*.

C: Our intuitions that moral deference is problematic *per se* do not show that moral deference is problematic *per se*.

The argument rests most heavily on P3, so I will focus on that. The literature has proposed two different factors to explain away the negative intuitions that we seemingly have about moral deference in itself. These factors concern aspects that have to do with the agents and their characters (as they are pre-deference).

The first factor is *ignorance* (Sliwa, 2012; Groll and Decker, 2014), and the story goes like this. If we look at some of the most disturbing cases of moral deference, we will see that the negative intuitions they give rise to are actually responses to the features of the agent who is deferring. Take the following examples:

Suit: Sam is standing at the shore of a lake when he sees a child beginning to drown. He believes that saving the child would be a good thing to do but it would involve ruining his new expensive suit. He cannot decide what to do and there is no one else at the lake, so he decides to call a friend whom he takes to be reliable. His friend tells him that he should save the child, and he believes him and saves the child. (Sliwa, 2012, p. 185)

Cats: Suppose Micah comes to believe that burning cats for fun is wrong simply because Eleanor has told him it is. Crucially, Micah's basis for his belief is not that Eleanor convinced him via argument, but simply that Eleanor told him that it is wrong. In other words, Micah came to his moral belief on the basis of testimony. (Groll and Decker, 2014, p. 54)

The optimists, together with the pessimists, would say that such cases of deference trigger negative intuitions and are indeed problematic. But, the

optimists add, what makes them problematic is the moral ignorance of the agent. That is the source of our negative intuitions.

The simple version of the ignorance argument (Sliwa, 2012, pp. 184-186) says that the issue here lies with the protagonists, who display an unusual level of moral ignorance. People who need to defer about whether they should save the drowning child when the cost is insignificant, or about whether torturing for fun is wrong, have very strange moral psychologies and dubious characters. That is, they are not aware of basic moral truths and have no idea how obvious reasons weigh up. One can wonder whether they can even be considered to be moral agents in the first place. This kind of ignorance in an adult moral agent just is very troubling. If we encounter someone who needs to defer about such issues, it seems that it is not the fact that they are deferring that makes the case problematic, but the initial ignorance itself.

The more complex version of the ignorance argument is based on the idea that there is a class of basic moral judgements that constitutes *normal knowledge*, which agents should generally have in virtue of being well-functioning humans (Groll and Decker, 2014, pp. 64-68). This normal knowledge (which can be knowledge-how or knowledge-that) consists of certain things we are expected to know and to do in order to be able to navigate our lives and the world. Regarding morality, this includes knowledge that torturing cats for fun is wrong or that we should save a person who is suffering at least when there is minimal cost to us. Not having this type of knowledge-that, and especially knowledge-how, like Sam and Micah, and needing to defer about it, shows a certain dysfunction in the moral agent. And *this* triggers negative intuitions. Whatever extends beyond knowledge which can be considered normal, however, can unproblematically be the object of deference. For example, in certain situation it is not so obvious how to weigh the relevant moral reasons, and then it would be fine to defer. Deferring about non-normal knowledge, the kind of knowledge that is not necessary for us as well-functioning humans, is not intuitively negative in character, these optimists hold.

Moreover, the ignorance argument extends beyond morality. This explanation can be applied beyond the moral domain to basically any area in which humans are supposed to have some minimal knowledge in order to be able to go about in the world; for example, knowledge of simple logic or basic

perceptual knowledge count as normal knowledge. Taking testimony about all such items of knowledge seems equally suspicious, while deferring about complicated mathematical equations does not. Thus, according to optimists, since the intuitively problematic character of these cases of deference extends to non-moral areas, and we do not conclude that there is anything wrong with non-moral deference from that, we should not do that in the moral case either (Groll and Decker, 2014, pp. 64-65).

Going further, there is another factor related to the agent's character that is said to explain away other cases of intuitively problematic moral deference. This explanation purports that sometimes people appeal to deference in order to achieve another aim rather than doing the right thing, such as to escape responsibility, to fit in, to impress someone else, and the like. Basically, they defer for some *ulterior motives* (Sliwa, 2012, pp. 188-189). Imagine Jack, who defers to Alanah about how much to help a person in need just because he wants to show her that she means a lot to him and that he would do whatever she says. Such a person clearly does not care, intrinsically, about what they should do regarding the person in need and what the right thing is. They are interested in something completely different, which has nothing to do with morality. So, it is this fact about their motivations that makes us suspicious of such cases of deference. Again, this works outside morality too: cases in which people who decide that others should make choices and judgements for them, for various dubious reasons, provoke the same reaction.⁹⁹

The problem with this negative strategy is that the best interpretation of pessimism can accommodate it. The negative strategy puts forward perfectly reasonable explanations of *certain* cases of moral deference that are problematic. It also offers a valuable lesson that a pessimist view should incorporate; namely, that we need to ask *why* a person defers. That can tell us a lot about the source of our negative intuitions and failing to take this into account will lead the pessimist astray. However, this optimist strategy nevertheless misses its target

⁹⁹ Sliwa (2012, pp. 186-187) mentions a third factor: *controversy*. There are some questions that people defer on that are controversial: rational, well-informed people still argue about them, and no consensus has been reached (e.g. the moral permissibility of eating meat). This is what makes such instances of deference problematic. This is, however, an epistemic factor. For clarity and simplicity, I will bracket this discussion here to avoid introducing variables that are extraneous to my purposes.

because it seems to misinterpret what pessimism is or should be.¹⁰⁰ A good version of moral deference pessimism would not be incompatible with this negative strategy of the optimist. Here is why.

Firstly, no pessimist needs to deny that some cases of moral deference are problematic due to some other reasons that have nothing to do with deference, be it ignorance or ulterior motives. People like Micah and Sam have some serious moral shortcomings and their cases can indeed be explained away by appealing to their moral ignorance. And no doubt some people do defer motivated by dubious ulterior motives. Not all instances of deference which trigger negative intuitions are bad for the same reason. There is no principled reason why pessimism should deny any of this. But such cases (that can be explained away) do not exhaust the range of problematic cases of moral deference.

Secondly, recall that pessimism says that for any case C in a certain range R (to be specified), moral deference is always *pro tanto* bad *per se* in C, even though it may be all-things-considered better to defer in C (and in many other cases, too). The definition is itself limited to deference that is problematic *per se*, i.e. not the kind of deference that optimists investigate in the negative strategy.

If the optimists want to flat out deny that there are any cases of moral deference that are problematic *per se*, then I think they would just be wrong. We cannot explain away *all* the problematic cases of moral deference that exist. For example, there are cases of *recurrent moral deference*, i.e. deference which happens repeatedly, which are *pro tanto* problematic *per se*. Imagine someone who, although not morally ignorant, morally flawed in any other way, defers on more than one occasion; let us say they do it quite a bit, sometimes because their own judgement is clouded, other times because they are uncertain and do not know what to do. Now perhaps each particular instance of deference does not seem problematic on its own (and as such it controls for the extraneous factors which preoccupies the optimist). However, I think, intuitively and reasonably, all of

¹⁰⁰ Sliwa (2012, p. 176) takes pessimism to operate under the following principle: “NO TESTIMONY: For a mature moral agent, there is something wrong with relying on testimony for one’s moral beliefs even if one knows one’s source to be reliable and trustworthy.” However, this is too general, vague, and simplistic a view of pessimism. I showed in chapter I that many more qualifications need to be made.

them together are indeed problematic.¹⁰¹ And the optimist negative strategy cannot account for the problematic character of such cases of recurrent moral deference. I thus conclude that the negative strategy of the optimists does not succeed.

2. The Positive Strategy

The main thought behind the optimist positive strategy seems to be the following: moral deference should be considered vindicated because it is not always bad. In some cases, it is good, but it can also be better, or even morally required, to defer rather than to figure things out on one's own.¹⁰² The views that fall into this category are more or less strong with respect to how good, or what kind of good, they take moral deference to be; more precisely, they differ in their stance on whether deference can be considered a second-best or "first-best" (as McShane, 2018, p. 640, puts it). I will present them in a scalar manner, starting with the weaker positions where deference is seen as a way of correcting our deficits and helping us get answers, where there are no other means. I will then end with the stronger views, where deference is said to bring about some distinctive goods that deliberating alone would not. I will discuss what I call *the coping argument* (Sliwa, 2012; Driver, 2015), *the no significant costs argument* (Enoch, 2014), *the epistemic injustice argument* (Wiland, 2017), and *the close relationships argument* (McShane, 2018). I will describe them in turn and then argue that they also fail to establish that we should always look on the bright side of deference.

2.1. The Coping Argument

According to *the coping argument*, moral deference is a way of managing, or compensating for, the different limitations and deficits that we have, intrinsically or due to our circumstances. This makes moral deference good and

¹⁰¹ Perhaps even the examples that the other pessimists have used could be improved to control for the extraneous factors discussed here and thus be shown to be problematic *per se*.

¹⁰² If this seems like a false contrast is because it is. The optimists take pessimism to be a lot stronger than it needs to be, so their arguments, as I will show, miss the point.

such cases of moral deference unproblematic. For example, recall the following example:

TRIP: Anna is a journalist who is preparing to go on a reporting trip to a dangerous and conflict-ridden area. She has to tell her family that she will be away but she really doesn't know what to tell them. If she tells them where exactly she's going and why, they will be extremely worried. On the other hand, she worries that by evading the questions she would be lying. She goes back and forth but cannot decide what the right thing to do is and eventually decides to ask a friend whose judgement she trusts. (Sliwa, 2012, p. 179)

Instances of moral deference like this one, it is argued (Sliwa, 2012), are not problematic because they can help us do the right thing when we cannot do that on our own. Sometimes we are simply uncertain of what the right thing to do is. But we can be lucky enough to have around us people who are in a better epistemic position and are able to figure out what the answer to our moral question is. They can be in a better epistemic position than us because it can happen that our judgement is clouded by prejudices or self-interest. When we realise that, we should appeal to the judgement of someone who is not similarly affected by these misleading factors. On different occasions, we may just think that other people have more developed moral sensitivities and, because of that, they are better at making certain moral distinctions. So, moral deference is good:

because we're in many respects creatures who make mistakes, who get distracted, who are susceptible to biases, who have limited abilities of discrimination in some areas and who are, moreover, well aware of all that. Moral advice allows us to do the right thing despite all these limitations by tapping into the cognitive resources of our peers. (Sliwa, 2012, p. 182).¹⁰³

¹⁰³ To be precise, Sliwa's argument is that relying on moral *advice* can sometimes be unproblematic, and that moral advice and moral testimony are not fundamentally different, and thus taking moral testimony is unproblematic as well. I briefly discussed the potential difference between taking advice and deferring in chapter I (p. 22) and decided to put this issue aside. For the purposes of this chapter, we need not be distracted by Sliwa's argument for why advice is the same as testimony, and so we can take them as interchangeable.

Moreover, precisely because of our deficits and our need to deal with them, moral deference can even be considered a *virtue*. Driver (2015) argues, first of all, that although it would be better to have no limitations such that we would not need to rely on anyone else's judgement (which already carves some room for pessimism),¹⁰⁴ being virtuous *requires* that we defer sometimes (given that we do have limitations). That is, we have imperfections, and part of what it is to be morally conscientious, responsible, and virtuous is to defer when we cannot make the right choice alone. This is because in those situations we need to be preoccupied with doing the right thing rather than with cultivating our moral worth and virtue. And doing the right thing sometimes requires that we defer to those who know what that is. So, if being a responsible moral agent sometimes implies deference, then, in this sense, moral deference is not incompatible with virtue (*pace* Howell 2014).

But, secondly, we can go further: deferring can be seen as virtuous in itself. More precisely, deferring about moral issues might be a *coping virtue*. A coping virtue is one that is beneficial to us precisely because of the (epistemic, moral, temperamental etc.) shortcomings we have. Deferring to others who know better is a way of coping and compensating for those shortcomings:

We will not be able to avoid all limitations, and, we will need deference as part of our decision-making tool kit, a quality that if responsibly deployed is part of an excellent human character. (Driver, 2015, p. 38)

2.2. The No Significant Costs Argument

The no significant costs argument is similar to the coping argument in that it portrays moral deference as a way of responding to a limitation: moral uncertainty. In this sense, it is also a second-best view of moral deference, presenting it as a coping strategy.¹⁰⁵ It is stronger than the previous views,

¹⁰⁴ Sliwa does not explicitly subscribe to this view and I am not sure if she would. But we do not necessarily need this to make the optimist case and my pessimist response will also apply regardless.

¹⁰⁵ In fact, Enoch (2014) explicitly offers an explanation for the fishiness of forming judgements by way of moral deference based on its lacking emotional and moral achievement. This is also why, as discussed in chapter I, Enoch falls in the middle on the pessimist-optimist spectrum: he looks like an optimist but does not walk like an optimist all the way. That is, he aims to vindicate

however, in that it insists that moral deference can not only be unproblematic, but sometimes *morally required*.¹⁰⁶ To see this, imagine a situation where you have to take an action with serious moral consequences. You have an opinion about what you should do but you doubt yourself, and the available evidence suggests that deference is more likely to provide you with the right answer. Given that you cannot refrain from acting in this case, it seems that if you do not defer, that you are taking a risk of wronging some people.

This scenario is a generalisation of Enoch's (2014) example about him deferring to his reliable friend, Alon, about whether to vote to fund a war. Alon usually makes the right judgements about this sort of things — and the fictional Enoch ends up agreeing with him all the time. Making the wrong choice here would have serious consequences, so he needs to make a decision. Enoch argues that because there are significant costs if his fictional self gets things wrong by figuring things out on his own, he is morally required¹⁰⁷ to defer in this case.¹⁰⁸ Deference, he holds, when other things are equal, is the appropriate response to moral uncertainty:

In fact, when someone genuinely believes that another is more likely to be right on a moral question, moral deference is not just permissible, but rather is morally required. Refusing to defer in such circumstances would amount to unjustifiably accepting a higher risk of compromising others' morally protected interests, or of wronging them. And this is inconsistent with giving others' morally protected interests the right kind of place in one's practical reasoning. It would be wrong. (Enoch, 2014, pp. 243-244)

moral deference and simultaneously offers an explanation for our negative intuitions about moral deference.

¹⁰⁶ I take Enoch's argument to be different from Driver's mainly because the latter operates in the virtue framework and the former does not. Driver discusses what is required for virtue, and Enoch what is morally required *simpliciter*.

¹⁰⁷ The 'ought' in question here is subjective: "The question we are interested in is what it is reasonable for me to do given my ignorance, uncertainty, partial evidence." (Enoch, 2014, p. 234)

¹⁰⁸ It is worth noting that Enoch has in mind primarily what Howell (2014) calls *active deference*, namely acting, rather than forming beliefs, based on deference (2014, p. 243). It is not clear whether his argument adequately extends to moral beliefs formed on the basis of deference. I am also not certain he would want that, since he offers an explanation for why deferential moral *beliefs* are fishy.

In other words, the argument goes, there is a moral requirement to minimise the risk of wrongdoing, and if we can do so by deferring with no other relevant costs, then we ought to defer. And it seems that there is no suitable candidate to make the other relevant costs matter sufficiently (*pace* the authenticity argument in Skarsaune, 2016). That is, not even the loss of things such as autonomy or moral understanding can be considered a significant enough cost to outweigh the value of minimising the risk of wrongdoing. This means that, although these things may have intrinsic value, we still ought not to risk wrongdoing just for their sake or for the benefits their cultivation might bring (e.g. an increased ability of acting reliably in the future; moral understanding of delicate complex moral issues). Moreover, it would be self-defeating (and wrong) for one to risk doing wrong in order to nurture their, e.g., moral understanding, when moral understanding itself is supposed to help us do the right thing and help, not harm, others. That person would not exhibit much moral understanding to begin with then. So:

there is no plausible candidate for a value here that could outweigh the general requirement to minimise the risk of wrongdoing. This means that we are back to the no-other-significant-costs scenario. Deference is the way to go. Moral deference has been vindicated. (Enoch, 2014, p. 250)

2.3. The Epistemic Injustice Argument

Notice, however, that the no significant costs argument focuses on moral deference as a response to moral uncertainty. But what if someone is not in such a situation, but is fully able to make the right decision by thinking things through on their own? In other words, can deference be *better* than non-deference? We are moving now into the ‘first-best’ territory.

According to *the epistemic injustice argument* (Wiland, 2017), there are instances in which not accepting someone’s moral testimony can be a form of epistemic injustice, where people are hurt in their capacity as epistemic agents (knowers). One motivation behind not taking testimony sometimes might be, though perhaps not explicit and voluntary, that we are influenced by certain prejudices which make us see others as less credible. For example, a woman

might not be believed by others when offering some piece of (moral) knowledge because of some prejudices that people have about her gender. This is an instance of testimonial injustice. Accepting the (moral) testimony of that woman, by contrast, can help fight the testimonial injustice and give people the epistemic status they deserve.

As a result of testimonial injustice, not only do we close ourselves to some moral knowledge that we might have rightly gained, but we can also harm others (in their capacity as knowers) whose credibility is unfairly doubted. Such harms include showing disrespect for certain categories of people as knowers and not giving them their proper place in the human community; dehumanizing and humiliating them, which may prevent them from getting the same opportunities as the rest; and making them fail to develop intellectual virtues and lose confidence.

Victims of hermeneutical injustice suffer in a similar way. Hermeneutical injustice occurs when “someone has a significant area of their social experience obscured from understanding owing to prejudicial flaws in shared resources for social interpretation” (Fricker, 2007, p. 147). Basically, in this case, someone cannot make sense of, and communicate, a social experience they have because there are no available adequate hermeneutical resources for them to use to make it intelligible. This is owed to the systemic prejudices that have been entrenched in our social identities by a discriminatory society and institutional setup. For example, for a long time women did not have the concept of ‘sexual harassment’ to talk about a particular social experience that many of them were having. This prevented them from comprehending what they were going through and talking about it with others (Fricker, 2007, pp. 150-151). The gap in the hermeneutical resources available made these women victims of hermeneutical injustice.

Victims of hermeneutical injustice suffer, first of all, from not being able to make sense of their own experiences. But they also suffer because they are not understood by others, and so they feel cognitively isolated. One consequence of this is, according to Wiland (2017, pp. 70-71), (epistemic) alienation. This line of thought (inspired by Marx) claims that if people do not understand each other and do not relate to one another in specific ways, they fail to live well together and fail to thrive.

Thus, deferring to a victim of such injustices can bring about some distinctive moral goods: it is a way of remedying the harms of such forms of epistemic injustice and it creates epistemic solidarity. By taking moral testimony from these disadvantaged people, one is fighting against epistemic and hermeneutical injustice, and giving these discriminated groups respect and their fair place in the epistemic community. One also helps people not feel alienated and welcomes them into an epistemic solidarity in which we can rely more on one another and we can benefit from each other's knowledge. As Wiland (2017, p. 72) puts it:

Trusting others' views about how to live doesn't merely make our own views more accurate, but also brings us together in a way that perfects our common humanity. Epistemic solidarity is part and parcel of the human good. There is thus just no way to be fully human without taking on the testimony of others.

Moral deference thus brings about not only epistemic benefits, but also some specific moral values. This is why, it is argued, the asymmetry thesis is true, but in the opposite direction: sometimes it is better to defer when it comes to moral rather than non-moral issues.

2.4. The Close Relationships Argument

In the same direction, we have *the close relationships argument*. McShane (2018) also argues that deferring on moral matters can bring about a distinctive good, specifically in intimate relationships, such as close friendships.¹⁰⁹ Depending on others by way of taking their moral testimony has a non-remedial value because it is an expression of trust, which is a central element of the ideal of friendship.¹¹⁰

¹⁰⁹ McShane (2018, pp. 631-633) claims that she is unique in putting forward an argument that does not present moral deference as second-best. She is right about most optimists, but nevertheless I believe Wiland (2017) does the same thing as her.

¹¹⁰ McShane (2018, p. 631) briefly states that the same argument could be made regarding moral advice (which she takes to differ from testimony, in that it is action-directed rather than belief-directed, according to McShane, 2015).

The story is the following: we depend on our friends for a lot of things, many of which are very important for us. This sort of dependence seems significant (when it does not go to the extreme or become pathological), as we tend to think that it brings people closer together. Being dependent on a friend in this way shows that we trust them, as it indicates vulnerability on our part. This seems to be “a core, constitutive element of an ideal of friendship” McShane (2018, p. 636).

Now one important thing that we can rely on others for is moral knowledge, which is central to our moral agency. Being dependent on our friends for moral knowledge can be a particularly good way to illustrate the trust we hold in them. We might even say that:

the ideal of friendship would be undermined by friends failing to be disposed to depend on one another’s moral testimony. That is, there’s good reason to think that being disposed to depend on one another’s moral testimony is essential to the ideal of friendship. (McShane, 2018, p. 638)

To see this, imagine that a friend, or your wife, invites you to believe the moral judgements they are making — a negative moral judgement about another person while recommending that you do not get to know them better; a judgement about the sexism of an advertisement they saw etc. Imagine you also simply do not accept the invitation and you do not form that judgement. You refuse their testimony and perhaps you go to figure things out on your own, if you can. Does that seem like something that should happen between two people who are close to each other?¹¹¹ This kind of behaviour does not seem to have place in the ideal of close relationships because it shows a lack of trust and a lack of willingness to depend on them in truly important ways. So, McShane concludes, deferring to our close friends about moral matters is a central part of friendships and of cultivating close relationships, and as such it has non-instrumental value.

¹¹¹ McShane refers to recurrent behaviour rather than isolated instances of not deferring, as she wants to talk about a disposition to be dependent on others for your moral knowledge.

3. Responding to the Positive Strategy

Here we have it: the positive optimist strategy which purports to show that there are cases where it is good, better, or morally required to defer about moral matters rather than figuring things out on one's own. These views, as far as I am concerned, make perfect sense. So much sense that, again, there is no reason for pessimists to reject them; and they do not, and do not have to, do so. My response to the positive strategy is fundamentally the same as that to the negative strategy.

To start, it would be too strong if pessimists would deny that moral deference can be an appropriate response to moral uncertainty or one's intrinsic and circumstantial limitations. Even more so if there is a risk of wronging others. But they do not. It is true that some versions say that one should try and cultivate one's, for example, moral understanding as much as one can, but it is unlikely they would recommend it when that involves doing wrong to others. Such cases are the unavoidable cases that the pessimists concede as unproblematic (see chapter I, section 2.3.2.). Recall that there are two ways in which we can understand this 'unproblematic' but neither causes trouble for pessimists.

First, it could be that it means all-things-considered unproblematic. This is probably the most accurate interpretation of what Sliwa, Driver, and Enoch say — after all, deference is treated as second-best in their arguments. But since no pessimist is committed to saying that moral deference is *always* and *all-things-considered* problematic, then we have no disagreement. The pessimist can allow that, at least in unavoidable situations (of uncertainty or epistemic haziness), moral deference is all-things-considered unproblematic. But the fact that they are unavoidable, and thus all-things-considered unproblematic, does not mean that they are not, nevertheless, *pro tanto* problematic and bad. So the pessimist can still hold that negative claim about moral deference. Recall that this is how I described the best interpretation of pessimism: for any case C in a certain range R (here: unavoidable cases), moral deference is *pro tanto* bad in C, even though it may be all-things-considered better to defer in C (and in many other cases, too). What the optimists say does not explicitly deny this, nor does it give us any reason to deny it.

Second, it could be that moral deference is unproblematic *simpliciter* in the unavoidable cases. That is, there is no sense in which deference is problematic. Pessimists could take this route as well and still hold on to their pessimistic hat. They can do that by saying only unavoidable situations are unproblematic. But that does not exhaust the range of relevant cases: there are cases where one could have formed the judgement or acted on one's own but deferred instead, and those can indeed be *pro tanto* bad. More precisely, I believe that recurrent moral deference can be *pro tanto* bad in such cases. I will show in chapter IV why. The point here is that the claims of the optimists who propose the idea that moral deference can sometimes be good or morally required do not exclude the possibility that those cases cannot be at least *pro tanto* bad even if all-things-considered good.

Going further, pessimism is also compatible with the scenarios that Wiland and McShane consider. This is because the pessimist can accept that there are cases where deference is indeed *pro tanto* good (perhaps even better) in some respect (which is what these optimists' claims are as well), to the extent that it decreases risk of wrongdoing, shows respect for knowers, contributes to fixing some systemic wrongs, creates solidarity, consolidates intimate relationships. At the same time though, the pessimist can hold that it is also *pro tanto* bad (or worse) in other respect, to the extent that (insert your favourite moral deference pessimism argument). All that the optimist has shown is that we should not always all-things-considered think things through on our own rather than deferring: but this is something to which the pessimist was not, and does not have to be, committed to begin with.¹¹²

4. An Interesting Pessimism?

To summarise, the main point of my criticism of the existing versions of moral deference optimism is that they misinterpret moral deference pessimism

¹¹² Wiland (2017, pp. 62-63) considers a similar conciliatory response, but paints the pessimist as wanting to say that, even in the cases he discusses, it would still be *even better* to form the judgement without having to rely on others. But this takes pessimism to be stronger than it needs to be. All the pessimist needs to say is that to there is a sense (to be specified by each view) in which it is *pro tanto* bad, to some extent, that the person did not manage to figure things out alone and deferred. This does not mean that deferring was not also *pro tanto* good or even all-things-considered better.

or they do not address its best interpretation. So they end up fighting windmills. For one, saying that moral deference is not even *pro tanto* problematic *per se* is extensionally inadequate, as it does not explain away all the problematic cases of recurrent deference. Secondly, claiming that moral deference can be *pro tanto* or all-things-considered good or better is compatible with the view that moral deference is always *pro tanto* bad (in certain cases).

However, a worry arises here: if moral deference optimism is compatible with moral deference pessimism, do we even have a conflict left? In the light of this agreement, what is the debate about then? I think the conflict still exists. If we take optimism to be the denial of pessimism, then its main claim should be that it is not the case that, in a certain range of cases (those specified by our pessimist views), moral deference (*per se*) is always *pro tanto* bad. Unfortunately, the existing optimist views are silent on pessimism interpreted like this. They do not explicitly consider the possibility of an ‘always *pro tanto* bad’ claim¹¹³ and do not formulate any arguments against it. Nevertheless, it would not be incoherent for them to deny this claim. The spirit, even though not the letter, of optimism is certainly so inclined. Now, of course, I have argued that the existing versions of pessimism fail and, if I am right, then their ‘always *pro tanto* bad’ claim with respect to individual cases does not stand up to scrutiny. Still, optimists cannot breathe easy. They need to deal with my version of moral deference pessimism, which I take to be the best version of pessimism: recurrent moral deference is always *pro tanto* bad. The optimist needs to put an argument against this view if they want to hold on to their stance.

But they might also go another way and challenge the relevance of such a view altogether. Are there any good reasons to understand pessimism in this way? Reisner and Van Weelden (2015, p. 442) worry about this so-called *pro tanto* problem: at best, pessimism can only put forward *pro tanto* arguments, but a *sans phrase* or all-things-considered pessimism does not seem plausible. This is a problem, they claim, because it is not very interesting. The authors do not explain exactly why it is not an interesting view. But it is suggested that such a position is not strong enough and it does not tell us much about moral deference — there might be plenty *pro tanto* reasons not to defer and probably none of

¹¹³ Although perhaps we can interpret them as implicitly arguing against such a claim when they argue against individual versions of moral deference pessimism.

them will be sufficient to establish that we should choose non-deference over deference *sans phrase*.

I agree that an always all-things-considered evaluative or deontic pessimism is out of the question.¹¹⁴ I venture to say that no one would want to argue for that, as it seems too strong. But is this really a problem? For one, I do not think anyone would want to embrace an all-thing-considered, *sans phrase optimism* about moral deference. But that does not seem to bother anyone. Nor should it — but then why worry about not having an all-things-considered, *sans phrase pessimism*? I believe that this gives us good reason to understand pessimism in this weaker way.

Secondly, and more importantly, the debate is about whether our negative intuitions about moral deference are on to something, whether they are warranted. Even if it turns out that the explanation for why they are warranted is not an all-things-considered one, the story still stands; we have an answer for what many people took to be an interesting question. If we think the question is valid and valuable, then why not accept a *pro tanto* answer, if this is the most plausible one? *Pro tanto* reasons still need to be taken into consideration.¹¹⁵

Sure, there are further questions about when the *pro tanto* value of non-deference — however that is specified — outweighs the possible *pro tanto* value of deference, i.e. how to weigh up the different *pro tanto* claims that can be made in certain cases where deference is an option. But how to weigh these up is a general issue for both optimists and pessimists, and the distinctive goal of pessimism is to tell a story about the problematic character of deference: when it is problematic and why. But we cannot ask that it tells us exactly what to do.

5. Appendix: Optimism outside Moral Deference Optimism

To end, I want to briefly discuss another potential challenge to the pessimist; a challenge that comes from outside the moral deference debate.¹¹⁶

¹¹⁴ But even this does not necessarily mean that it is *impossible* for the value of figuring things out by oneself to ever outweigh any other value. Although this does not warrant a stronger pessimist claim, it is very possible that, *in certain cases*, it might indeed outweigh other values.

¹¹⁵ This is also why I think that, with all the concessions, I am still a moral deference pessimist.

¹¹⁶ I thank Matthew Smith for raising this suggestion.

The debate on authority, specifically in the political and the legal domains,¹¹⁷ takes active deference — that is, acting, rather than forming beliefs, on the basis of deference — to be unproblematic.¹¹⁸ More specifically, it is often argued that when there is a legitimate authority in place, we have reasons, and possibly an obligation, to defer to it.¹¹⁹ Sometimes, the directives of the authority will have moral content. This implies then that we have a duty to morally defer to legitimate authorities.

To get a better understanding of the view, consider one way of conceiving of legitimate authority: Raz's instrumentalist view (Raz, 1986).¹²⁰ On this account,

the normal way to establish that a person has authority over another person involves showing that the alleged subject is likely better to comply with reasons which apply to him (other than the alleged authoritative directives) if he accepts the directives of the alleged authority as authoritatively binding and tried to follow them, rather than by trying to follow the reasons which apply to him directly. (Raz, 1986, p. 54)

In other words, an authority is legitimate if obeying it will increase one's conformity with the reasons that exist. That is, one will act in accordance with them more often if one defers rather than if one tries each time to think about things unaided by anyone. This is *the normal justification thesis*.

The normal justification thesis is complemented by a couple of other claims. One is *the dependence thesis*, according to which the legitimate authority is supposed to issue directives that are based on the independent and relevant reasons that already apply to the subjects of those directives. This means that, for example, if the authority issues a new directive, say, the introduction of a new traffic sign, it will do so for the reasons that warrant that action. Namely,

¹¹⁷ There are certain differences between legal and political authorities, but it will not matter for my purposes. For a discussion on how Raz applies his view on political authority to the legal setting see Ehrenberg (2011).

¹¹⁸ The exception is the philosophical anarchist position; see Shapiro's (2002) review of this type of argument.

¹¹⁹ Not on all accounts of authority, however. See Christiano (2013, sections 1.2-1.3).

¹²⁰ I am aware of the criticism that Raz's views have faced. For overviews see Ehrenberg (2011); Edmundson (1993); Whiting (2016); Perry (1989). However, since my focus is something else, I will not address them. Instead, I use his position only as a prop to ask further questions about moral deference.

that it will lead to more safety on the road for both drivers and pedestrians, it will help the community overall etc. The authority will not issue the directive for reasons such as: the people in charge want to earn some extra money and they could do so this way, the people in charge want to seem like they are contributing to the life of the community, and the like. Directives based on such reasons would not be legitimate because they would not meet the dependence thesis.

There is also *the independence thesis*, which says that the matters that an authoritative directive targets should be such that it is not more important for the subjects to think through them on their own rather than simply conform to reason (Raz, 2006). This basically restricts the jurisdiction of authorities to those situations where acting correctly, in conformity to reason, is what matters most. One example regards directives concerning the use of medicinal drugs; decisions about safety of pharmaceuticals is not the kind of personal thing that we should do by ourselves, unlike, e.g., deciding whether to take a certain medical treatment. It is better to leave decisions about general safety of medicinal drugs to the relevant authorities and we should just defer to them.¹²¹

Raz's view has also been called *the service conception of authority* as he considers the role of political authority to be that of a mediator between people and reasons. What the authority does is to take into account all the reasons that are relevant to its subjects in a situation and be sure that it issues directives that conform to the balance of those reasons that apply. For the authoritative directive to have the purported force and achieve its goal, the authority also has the power to pre-empt some reasons. *The pre-emptive thesis* (Raz, 1986, p. 46) states that a directive issued by the authority is not to be added to the relevant reasons that the subject is considering when deciding how to act. Instead, it excludes and replaces the reasons that the subject has already taken into account. For example, to obey a co-ordination related directive, such as driving with no more than 30 miles per hour on restricted roads, is not to do that because it is the law *and* it is safer. The fact that it is the law excludes our acting for safety reasons because

¹²¹ I think this shows that the duty of obedience that Raz discusses should not be understood as an absolute duty, but more like a *pro tanto* one, with respect to its satisfying the normal justification condition. The Razian account is also piece-meal: it admits that while an authoritative directive might apply to me, it might not apply to someone else, if the normal justification thesis is not met or if it's more important to act on one's own; or it might apply to me now, but will not in another situation (Ehrenberg, 2011, pp. 886-887).

the law is already based on those reasons; otherwise, we would be double-counting the reasons.

Instead, authority creates a new kind of reasons for which we act, namely *content-independent reasons*. A content-independent reason is a reason that calls for action not in virtue of the goodness of what that action prescribes, but because of something that is indifferent to its merits — here: because it came from an authoritative source. For example, an authoritative directive is a content-independent reason to drive under 30 miles per hour on restricted roads. We have a duty to drive like that not directly because of the merits of such action, but because the authority demanded us to do that. If a general gives a command to a group of soldiers to intervene in a conflict, that order is a content-independent reason for them to intervene. The soldiers do not act for the reasons that warrant the intervention, but solely on the general's command. Such reasons are to be contrasted with *content-dependent reasons*, which are reasons to do something because of the content of that action. For example, I act on content-dependent reasons if I donate to charity because I think the cause is worthwhile and it will provide some much-needed help to others (and not because someone ordered me to do so).¹²²

So, the legitimacy of authority and the content-independent reasons that it creates ground the duty of deference. More specifically, since legal and political authorities are practical authorities, the duties they impose are practical duties, i.e. duties to act in accordance with them. Although they might give indirect reasons for belief, they mainly give reasons for action. The subjects of such an authority will have a duty to act on the authoritative directives issued, but judge and deliberate as they please (Raz, 1986, p. 42). However, this still implies that, if the commands of the authority have moral content and take into account moral reasons, then we have a duty to defer about moral matters. For example, a political authority may put in place a certain taxation system. Or a legal authority may put forward a law of abortion. If these are legitimate authorities, then we ought to defer to them, and do as they say: comply with the taxation system and with the abortion law.

¹²²A similar distinction is the one between *first-order* and *second-order reasons*, where the former refer to the reasons that we have for or against an action and the latter to reasons to act or to refrain from acting for a certain reason, Raz (1999, p. 39).

This is a rich and fruitful topic in political and legal philosophy, but I think that transferred to the moral deference debate it amounts to no more than another version of the coping argument. That is, sometimes we will have good reason to think that we are not able to weigh up the moral reasons in a correct way or that someone else is better than us at that. That gives reason to defer, especially if we need to make a decision or act. And this is precisely how the normal justification thesis grounds the duty to defer in action — we have this duty *only if* we are less likely to do better when making the necessary moral calculations of the relevant reasons on our own.¹²³ But, again, the pessimist need not deny there are circumstances in which it is all-things-considered good to defer. Moreover, it is still open to the pessimist to say it would be *pro tanto* better if they could comply with the reasons that apply to them by reflecting on their own (for various reasons — insert your preferred pessimist argument here).¹²⁴

Moreover, perhaps the pessimists can have (at least some of) their cake and eat it too: recall that this view only asks of us to defer in action, it means that it is open to us to reflect on the authoritative directives as much as we want. This way we can perhaps still gain moral understanding or whatever it is the pessimist wants to hold on to; all of this whilst we are sure we do the right thing and not risk wronging anyone, as we listen to the directives that are likely to lead us on the good path (which will hopefully overlap with the answers we get if we figure things out on our own).

6. Concluding Remarks

In this chapter I argued that moral deference optimism is not successful in eliminating moral deference pessimism. Optimism, as it stands, is actually compatible with the most plausible version of pessimism. One might worry that what I call the most plausible version of pessimism is too weak. I disagree and

¹²³ It is an open question how we can establish this. In our non-ideal world, it might be very difficult to figure out whether you are more or less likely to comply with the moral reasons that apply to you by doing something as you yourself see fit or by following an authoritative command. This is simply because we might not have the necessary epistemic facts about the person in authority to make this judgement.

¹²⁴ Additionally, since this duty to defer is not absolute itself, depending on the details of the situation, it may still be weighed against other *pro tanto* values that we have (see the independence thesis).

in the next chapter I will present my own version of pessimism, one that I hope will convince the worrier that a *pro tanto* pessimism can be interesting and informative.

Chapter IV: The Real Problem with Moral Deference

In the previous three chapters I did an inventory of the current debate on moral deference. I argued that no existing account gives a satisfactory explanation of why moral deference is, or is not, problematic, and that we have yet to get to the heart of the problem. On the one hand, the existing moral pessimist views have shortcomings that prevent them from convincingly establishing the case against moral deference. On the other hand, the existing moral optimist views also fail to completely redeem moral deference. So, first, I proposed a different way of understanding moral deference pessimism. I hold that recurrent (moral) deference is always *pro tanto* bad, even if deferring is not *pro tanto* bad in individual cases, and even if some patterns of recurrent (moral) deference are all-things-considered better than not recurrently deferring in those ways. Now I will introduce my own, novel pessimist account, which will focus on recurrent cases of deference. In what follows I will show what makes *recurrent* moral deference *pro tanto* bad. I hold that my version of pessimism will show what the real problem with (moral) deference is. This account will withstand all the challenges I posed for the other pessimist views and will have certain independent advantages over them as well. This new variety of moral deference pessimism can also accommodate the existing optimist story and still vindicate our negative intuitions about moral deference.

My explanation of what makes recurrent moral deference problematic draws on its negative interference with our capacity for practical deliberation. My argument consists of two main claims. The first one, *the value thesis*, trades upon the significance of the capacity for practical deliberation in our lives, i.e. its instrumental and (extrinsic) final value. The second claim, *the interference thesis*, shows that recurrent moral deference interferes with the exercise and the development of the capacity for practical deliberation. I hold that the two theses, together with the idea that it is *pro tanto* bad to interfere with something of value, point to the conclusion that recurrent moral deference is *pro tanto* bad.

All by itself, this argument does not rule out the possibility that the exercise and development of the capacity for practical deliberation is not the only valuable thing with which recurrent moral deference interferes. For instance, perhaps it also interferes with the exercise and development of our

capacity for moral understanding. However, in chapter V I will argue that my specific view is better because it grounds a more satisfactory view about the asymmetry of deference.

1. Moral Deference and Practical Deliberation

Let me start with some scene-setting. To situate the problem, note when and how deference occurs. Here is one plausible take on how things go when someone wants to answer a moral question but does not manage to do it straight away. The agent is faced with the moral issue. They are not able to think of an answer instantly, as perhaps it is a complex issue. Say neither a heuristic (or any other fast and intuitive response), nor a habitual response, is available here. For example, think about the decision regarding whether or not to have children in this politically unstable world, which is also threatened by climate change. This is a hard question, which many people could not answer immediately.

What could a typical agent do in such a situation? One thing — plausibly what most people do — would be to do some *practical deliberation* to try and thoughtfully figure out what to do. However, there is at least one other way of getting an answer: they can defer to another person. Instead of deliberating, the agent could just ask someone virtuous and reliable what to do. Deference could then either replace the process of deliberation completely (i.e. one defers without reflecting at all on the question) or intervene at different stages, by supplanting one or more of the elements involved in deliberation. For example, one can defer about what considerations are relevant to the question, how these weigh up, or what matters most. Instead of figuring this out by oneself, one just uses a piece of testimony. I will come back to this later, but for now it suffices to say that all of this means that when confronted with a complex question, the agent has a choice at least between deference and deliberation. If one chooses the former, the typical process of practical deliberation will not occur as it normally would, as one skips one or more of its different steps by taking the shortcut of deference. This point, as I show below, is what helps us to get to the heart of what truly is problematic about deference.

Does it matter which way the agent chooses to solve their moral question? And, if yes, why? In answering these questions, I propose a new

version of moral deference pessimism, elaborated in the argument I defend in the following sections:

P1: It is *pro tanto* bad to interfere with something of value.¹²⁵

P2: The exercise and the development of one's capacity for practical deliberation are (both instrumentally and extrinsically finally) valuable. (*the value thesis*)

C1: It is *pro tanto bad* to have one's exercise and development of the capacity for practical deliberation interfered with.

P3: Recurrent moral deference interferes with the exercise and the development of one's capacity for practical deliberation. (*the interference thesis*)

C2: Recurrent moral deference is *pro tanto* bad.

In what follows, I take P1 as an assumption,¹²⁶ and I will first argue for *the value thesis* and then for the *interference thesis*. But before that, I elaborate on the notion of practical deliberation I have in mind. We use our capacity for practical deliberation to engage in instances of practical deliberation, which are a kind of reflective processes of trying to figure out what to do. This typically involves thinking about what considerations are relevant to the question, how they weigh up, or what the right answer could be. The deliberative agent perhaps tries to get more information about the available options and their consequences. They imagine how these possible courses of actions might unfold, what they would imply, and what they would mean to their life. The agent might go back and forth between the different options, trying to figure out the moral relevance of different aspects that seem to play out in the imagined scenarios, in accordance with their ends and goals. They could think about what the relevant values specifically mean to them, how they apply in the situation, and perhaps how they stand in relation to each other, especially if they diverge. This process usually

¹²⁵ By interfering with a value I mean here something like not respecting that value (or preventing its manifestation in situations where it is fitting). When I argue for the interference thesis it will become clear in what sense and exactly how recurrent deference does not respect the value of practical deliberation (i.e. precluding it from its development or contributing to its decline).

¹²⁶ P1 is not intended to be a strong general normative principle. Rather, I take it that it is plausible that normally, or at least in some cases, it is bad to interfere with something that is valuable, to the extent that the thing is valuable. Some things are good, important, and have value for us and, in that respect, it is at least sometimes bad to interfere with such things. For the purposes of my argument, I ask the reader to grant me that practical deliberation is one of those things, such that P1 applies here.

ends with the agent getting an answer, making a choice, a decision, or forming an intention. This is what I take a process of practical deliberation to be. The capacity for practical deliberation is what enables us to engage in such processes.

Practical deliberation, as I think of it, is not limited to instrumental reasoning, but can be about ends and not just about means (Kolnai, 1962; Richardson, 1997). Moreover, it includes things like choosing the inputs to deliberation (Arpaly and Schroeder, 2012), weighing up these inputs, using our imagination, and finding constitutive solutions to our problems (Williams, 1981). Practical deliberation need not take place at the forefront of the deliberator's mind: they do not necessarily think about the stages of deliberation they go through, conceptualize or articulate everything in a clear or theoretically sophisticated manner. It is a deliberative 'calling things to mind' type of process, but the agent need not be thinking about it in these terms.¹²⁷

Moreover, practical deliberation has another important feature: it is a way of tracking and responding to reasons.¹²⁸ This is probably why we generally decide what to do by deliberating rather than flipping a coin or guessing. We want our actions to make sense, we normally do things for reasons — something that flipping a coin does not typically provide (except, perhaps, in some special cases). Recall the example about whether to have children in today's unstable world. You would probably not want to make this decision by simply *guessing* what you should do. You would want to think about the pros and cons, and make a decision based on reasons.

Practical deliberation can help us get there.¹²⁹ When we reflect on the considerations that are relevant to our practical conundrums, when we imagine what we could do and how it would go, we are trying to find and grasp the reasons that apply to us. Then we react to, and use, those reasons to make decisions, formulate intentions, or act. This is what we typically take the role of

¹²⁷ However, this automatic thinking might be influenced by our conscious practical deliberation in different ways (see, e.g., Arpaly and Schroeder, 2012, p. 234; Kahneman, 2003, p. 710).

¹²⁸ Non-deliberative types of thinking and acting (fast or automatic thinking; habit) also have this feature, but they are not available in the cases I am interested in.

¹²⁹ Many philosophers strongly support the idea of a connection between practical reasoning and reasons-responsiveness, e.g. see Kauppinen's (2018) overview. Although they broadly overlap, I focus on practical deliberation instead of practical reasoning because I am interested in the cognitive phenomenon that encompasses more than practical inferences or syllogisms (i.e. that includes, among other things, how we choose the inputs to these inferences), which is how many philosophers discuss practical reasoning; e.g. Audi, 2004).

practical deliberation to be (Tiberius, 2013). For my purposes, there is no need to be very specific about what reasons-responsiveness is. However, I will say that, following Fischer and Ravizza (1998), I am referring to the process of recognizing and reacting to reasons. The former involves grasping the reasons for or against a specific action that exist and apply to us, and the latter implies translating those reasons into choices and actions. The recognition or receptivity component concerns the agent's recognizing particular considerations as (sufficient) reasons for an action. For example, I recognise that having to finish my paper is a reason not to go with my friends to see the ballet. Or I recognize that the pain that animals go through to sustain the food industry is a reason not to eat meat. The reactivity component is about the agent using those reasons to make choices, form intentions, or act. For example, I decide not to go to the ballet and tell my friends that I need to finish my paper instead. Or I make the decision to stop eating meat and I just stop eating meat.

To end this section, it is worth noting that my argument will not require taking practical deliberation as the only or the best way to think about things; not at all. There is also the aforementioned fast, automatic, non-deliberative thinking which we often use and which can be just as good, if not better in certain circumstances (Kahneman, 2003, p. 698; Railton, 2004, pp. 187-188; Raz, 1999). For example, by using the quick heuristic type of thinking we can solve simple maths problems; make certain associations; understand simple sentences; have certain reactions to certain things (e.g. disgust when seeing roadkill) etc. We are the kind of creatures who use both kinds of thinking, the automatic and the deliberative one. Both are important and have their place in our lives, and I do not rank them in order of importance or usefulness here. But, for my purposes, conscious practical deliberation is the relevant focus because that is the type of cognitive process that is at the centre of the cases of deference that I discuss, since the automatic type of thinking is simply not available in said cases.

1.1. The Value Thesis

Now that we have a grasp of what practical deliberation is, I will provide support for the main premises of my argument. I start by arguing for *the value thesis*: the exercise and the development of one's capacity for practical

deliberation are both instrumentally and extrinsically finally valuable. I primarily have in mind here the well-functioning of the capacity, i.e. cases where one exercises and develops it in the right direction, conducive to good deliberation. Of course, one can also do that in a bad way, conducive to inadequate, incomplete, or bad deliberation. Nevertheless, most of what I say below applies to such practical deliberation too, insofar as it has the potential to develop and turn into good practical deliberation.

Begin with some common-sense intuitions about practical deliberation in order to see the kind of place it has in our lives. Trivially, the capacity for practical deliberation is what enables us to engage in instances of practical deliberation in order to figure out what to do. It can help us answer our moral questions, but it is not limited to that. By definition, it covers all practical questions. Of course, as said above, practical deliberation is not always the best way to figure out what to do, as we also use automatic thinking or habit. Nevertheless, it is hard to imagine creatures like us not using it at all. After all, it seems that we typically achieve automatic thinking or the habit of acting intentionally but unreflectively only after we have gone through some instances of practical deliberation. The competent driver does not deliberate about what to do in a traffic jam anymore because she has mastered that skill by repeatedly deliberating, deciding what to do, and learning from that. The teacher quickly knows how to deal with the misbehaving student because he has experience with that kind of situation and previously thought about how it should be handled best. We can usually track automatic or habitual action to some past use of the capacity for practical deliberation. Moreover, when we are faced with a complicated issue or a conflict of values, heuristics or habit cannot always help us. Then we need to reflect and to deliberate.

So, firstly, given its important role in decision-making, it is highly plausible that the capacity for practical deliberation is *instrumentally valuable*.¹³⁰ I take practical deliberation to be both causally and constitutively valuable. That is, valuable because it is a means to obtaining something else that has value or a constitutive part of something else that has value. First of all, since practical deliberation is, above all, a way of deciding what to do, it is causally

¹³⁰ I will assume here that there is such a thing as instrumental value. This is widely accepted in value theory, but not everyone agrees (e.g. Rønnow-Rasmussen, 2002).

instrumentally valuable because it helps us navigate the world and answer our practical questions, in a reasons-responsive way. When we are in a complicated situation that we cannot solve instantly, practical deliberation usually represents our best chance of arriving at the right answer. Knowing how to look for reasons, how to recognize which considerations are relevant and how they weigh up, how to choose between different possible courses of actions, and so on, is a very important tool for us. From the most trivial to our most significant decisions, practical deliberation helps us make choices and act in response to the reasons that apply to us. This is why it is so difficult to imagine ourselves without it. Indeed, consider Enoch (2011, p. 70) who argues that not only is deliberation important, but it is a non-optional project for us, “partly because we are essentially deliberative creatures”.

Secondly, practical deliberation is also causally instrumentally valuable to us because it makes us act in a reliable and safe way. This is particularly important given the kind of world we live in. We might not always have a choice between deliberation and deference, as the latter will not be an option all the time. We might not be next to someone more knowledgeable or reliable whenever we need testimony. So, it seems important to be able to deliberate and be reasons-responsive on our own if we want to answer our practical questions, solve our moral conflicts, and generally become reliable in doing the right thing. Practical deliberation is valuable because it is not based on luck, but rather on reliably responding to reasons and safely reaching practical conclusions. By ‘reliable’ I mean that it is conducive to doing the right thing or likely to yield in correct results or decisions, across a rather wide range of situations. By ‘safe’ I mean that it would have not been easily wrong or misled. This is also what helps us ensure that our future selves’ decisions will also be the right ones.

Finally, it is plausible that this capacity also has constitutive instrumental value, as it seems to be a part of many of the other things that we find valuable. For example, we may think that being able to figure things out on one’s own is partly constitutive of autonomy, whether we think of that in terms of self-governance or good governance (Smith, 2004). In order to make decisions for oneself and be autonomous, one will need to use practical deliberation, at least to some extent. Similarly, we could make a case for saying that practical

deliberation is typically constitutive of leading an authentic life.¹³¹ If we want our beliefs and choices to express who we really are we will need to be able to figure out what kind of values we hold and how to apply them in different situations, or how they weigh up against one another.¹³² Virtue also seems to require practical deliberation, at least some of the times. Applying different virtues in particular circumstances, thinking about what a virtuous exemplar would do, using one's practical wisdom — all involves exercising one's practical deliberation.¹³³

Thus I think that all these points strongly support the idea that our capacity for practical deliberation is instrumentally valuable, both causally and constitutively.¹³⁴ Insofar as we want to act in accordance with the reasons that apply to us, and insofar as practical deliberation is a particularly effective way of doing that, practical deliberation is instrumentally valuable. But I think we can go further and say that the capacity for practical deliberation is more than merely instrumentally valuable. Now for my purposes it is not *crucial* to have an exact classification of what kind of value that would be. My point gets off the ground as long as I can show that it is plausible that practical deliberation is *not only instrumentally valuable*.

Nevertheless, the way I think about it comes closest to what has been called *extrinsic final value*. Something is extrinsically finally valuable when it is valuable as an end, but not due to its intrinsic properties; rather, it is valuable “in virtue of an external property it possesses—in virtue of its relation to some external object or objects” (Rabinowicz and Rønnow-Rasmussen, 2015, p. 31).

¹³¹ Except perhaps for special cases where someone's authentic self is defined by a lack of ability to practically deliberate or be reasons-responsive.

¹³² See, for example, how some philosophers talk about integrity (which can be taken to be quite similar to authenticity). About the view of integrity as self-integration it is said: “A person is subject to many conflicting desires. If one simply acted at each moment out of the strongest current desire, with no deliberation or discrimination between more or less worthwhile desires, then one clearly acts without integrity” (Cox et al., 2018). The view of integrity as moral purpose, based on Halfon (1989), is described as deeming “integrity as centrally concerned with deliberation about how to live” (Cox et al., 2018).

¹³³ It may be possible that one can be virtuous without doing any practical deliberation at all. However, many virtue ethicists deem practical wisdom at least a constitutively part of virtue. And practical wisdom is just a more specific type of practical deliberation that operates in the virtue ethics framework. For example, according to Hursthouse (2001, p. 13): “Each of the virtues involves getting things right, for each involves phronesis, or practical wisdom, which is the ability to reason correctly about practical matters.”

¹³⁴ One can disagree with the idea that practical deliberation is constitutively instrumentally valuable in the way described. But even if that is true, it is reasonable that it will still be causally instrumentally valuable to those things, and that is enough to get my point off the ground.

For instance, to use Rabinowicz and Rønnow-Rasmussen's example, Abraham Lincoln's pen may have value for its own sake exactly because it belonged to an important figure (and not due to some intrinsic properties of the pen). I hold that practical deliberation has value as an end because of its relation to the lives of creatures like us. Creatures like us value the exercise and development of the capacity as an end because it is the procedural skeleton of a sort of life that is important to us. That is, a life where we get to set and specify our own goals and construct our normative outlook of the world.

To begin to see this, we can ask something like the following questions: would anything be lost if we conceived of practical deliberation as a mere means for something else of value? One answer is that if practical deliberation were merely instrumentally valuable, then it would not matter much if we replaced it with another means that gets us the same thing, namely answers to our practical questions. For example, we could invent a pill that, when ingested, just eliminates the need for practical deliberation as it generates in our minds the same answers practical deliberation would. Or, for that matter, we could simply defer to other people all the time. Would anything be lost? My intuition is that some of us, maybe most of us, would not take the pill. Moreover, we have already seen that many people have reported the intuition that at least some types of deference are indeed strange. I believe this is because something of value would indeed be lost if we simply replaced practical deliberation. We would be forgoing a kind of life that seems valuable to us, one which could only be constructed through the exercise of practical deliberation.¹³⁵ That is, a life where we get to make most of our own decisions and choose how to construct our ends and commitments, and shape who we are. This suggests, I take it, that this capacity is not merely instrumentally valuable.

The reason I take practical deliberation to be (extrinsically) finally valuable is this: the exercise of the capacity, as a procedure which involves reflecting on considerations, taking perspectives, weighing reasons, making choices etc., is a way through which we shape, define, and specify our values, ends, and normative commitments. As imperfect moral agents we do not have

¹³⁵ Some people might want the pill or, indeed, defer all the time. This thought experiment is not supposed to be decisive, but just a way to awake some intuitions about there being a possible reason to think practical deliberation could be more than instrumentally valuable.

all the right answers or a list of rules that we can always straightforwardly follow. Even if we did have a list of rules, we would still need to use our judgement on how to concretely apply them. We also often need to figure out how our values apply to different situations because we learn new things, we change our minds, we correct mistaken views, we adapt and respond differently to different situations.

For example, most of us generally value things like friendships, knowledge, love, good characters. But these values can be expressed in many ways and might mean different things to different people. Each of us will go through different experiences and there will be times when, to form beliefs and to act, we will have to specify exactly what a friendship is for us, what kind of knowledge we want, what love amounts to, or how we should act to build a good character. Moreover, this might change with time, with context, with the situation, with the information we gain, and so on. So, it seems that with many of the normative questions that give us pause, with many of the decisions we need to reflect on, we shape and give nuance to all the things that shape our worldview. And we normally do it (at least in part) by using practical deliberation.

Here I build upon on the arguments of the philosophers who hold that practical deliberation can be about our ends, and not (just) about the means we take to reach our ends (Kolnai, 1962; Richardson, 1997; Wiggins, 1976). They claim that practical deliberation will qualify the practically achievable specifications of our ends. It will change and reorder our concerns, and it will aim to solve conflicts of divergent goals:

Our ends are not all ready-made, awaiting their fulfilment when the proper means should have been found; they may come to life and harden into shape in fairly unexpected contexts; and their fixation involves to some extent, at times it may be a considerable extent, a revision, modification and reorientation of our preestablished structure of permanent or comparatively lasting ends - I would rather say, our *concerns* - itself. It is the choices, confrontations, inner dialogues, hesitations and new engagements implied in this process that *primo loco* constitute the field of deliberation. (Kolnai, 1962, pp. 205-206)

If practical deliberation does indeed do this, then I think it is plausible to that it is more than merely instrumentally valuable. Practical deliberation is not just a means to obtaining more specific goals and normative commitments. We could get that in other ways, e.g. the pill or constant deference. But something valuable would be lost then: a certain kind of life where *we* define who are and how we see the world by specifying our projects and aims, by setting our commitments, and particularising our values. Practical deliberation is more than just a means to living a certain kind of life that is valuable for creatures like us, but a way of shaping that kind of life itself.¹³⁶ This is why I think it is extrinsically finally valuable too.¹³⁷

1.2. The Interference Thesis

We have seen what the capacity for practical deliberation is and why it can be considered valuable. Now I want to explain how moral deference interferes with it. As discussed, it seems that individual instances of moral deference interrupt individual processes of practical deliberation. The typical process of practical deliberation just does not unfold as usual because it is, or some of its different steps are, replaced by deference. Take this example: I need to decide whether to lie to my friend about whether I am really interested in his new hobby, knitting. He wants to talk about it a lot, but I do not find it that riveting. Because he is a sensitive person, I know I will hurt his feelings if I tell him the truth and he will feel like he cannot share the things he loves with me, which I do not want to happen. So, I need to decide whether to lie about how interesting I find stories about knitting or to tell him the truth and risk hurting his feelings. Because the answer does not just come to me, I need to deliberate. More specifically, I need to think about what considerations are relevant, e.g. what dishonesty entails, his dislike of dishonesty and preference for honesty,

¹³⁶ If one does not think that this is important, not least just because humans generally do it in order to be able to coherently navigate their practical lives, then I am not sure anything I can say would change one's mind. So I will take that it is good to live such a life as my bottom turtle.

¹³⁷ I am aware of the ongoing debate in value theory about different types of values, and of the fact that not everyone uses the terms in the same way and does not agree with all the possible distinctions (Rabinowicz and Rønnow-Rasmussen, 2003; Tucker, 2019). What I want to say about the value practical deliberation is that it is more than instrumentally valuable, but not that it is intrinsically valuable.

what kindness entails and how it is best to specify it here, how honesty and kindness weigh up against each other; I imagine his reaction if he later found out that I lied versus his reaction to a sensitive confession that I hate knitting, and so on. Now instead of doing any of this, I can also defer. I can simply ask someone I consider reliable before I deliberate at all and just take their testimony. Or I can defer at any stage in the process of deliberation. I can defer about which considerations are relevant, how they weigh up, or about how to specify my values. In that sense, I would not go through practical deliberation as I would in a deference-free scenario.

When an agent appeals to deference repeatedly, we have multiple instances of practical deliberation that are circumvented in the way described above. If I defer recurrently, it seems that, at least in some sense, the capacity for practical deliberation itself is interfered with. More precisely, *the exercise* of the capacity for practical deliberation, as a temporally extended process, is interfered with. This seems true simply in virtue of our idea of what a capacity is: if we exercise a capacity to Φ , there will be some instances of Φ -ing. If we interfere with different instances of Φ -ing then we will interfere with the exercise of the capacity to Φ .

But I think there is more to this: besides interfering with the exercise of the capacity for practical deliberation, it is plausible that recurrent moral deference also negatively impacts its *development*. Start with some common-sense ideas about capacities. We tend to think that my capacity to, say, play the piano can improve if I practise playing the piano. If I want to develop this capacity, I will not just read about it or watch others play. Development seems to imply exercise. Intuitively, we think that the more I practise, the more my capacity develops. This is how I build up my musical sensitivity, how I learn to distinguish good notes from bad notes, how I improve my finger placement and speed of playing, and how I eventually become able to play difficult songs.

Many empirical studies support the idea that exercise is crucial for capacity development. According to *the deliberate practice account* of expert performance acquisition, to reach an expert level in exercising a capacity, one needs an average of ten years of deliberate practice.¹³⁸ Deliberate practice is

¹³⁸ The studies were done in music (Ericsson et al., 1993), chess (Charness et al., 2005), spelling (Duckworth et al., 2011), and typing (Keith and Ericsson, 2007).

defined as a set of structured activities undertaken with the aim of improving a capacity, and is distinguished from play and paid work. It requires continuous effort and dedication, it is not inherently enjoyable, and it needs to include adequate feedback from an instructor.

However, this research has its critics, with some psychologists suggesting that practice is not the whole story. They argue that other factors, such as genes, intelligence, or starting age, play a role in the level of development of a capacity (Hambrick et al., 2014; Meinz and Hambrick, 2010). Still, the authors do not deny the importance of deliberate practice: they deem it a necessary factor, although not sufficient for attaining expert level. But this is precisely the key point to take from these studies: deliberate practice is at least a *necessary* element in the development of a capacity to an expert level. Given this, if anyone wants to deny that practice is not needed to highly develop some capacities, the burden of proof is on them. Now it is true that the empirical research does not say anything about more ordinary levels of performance, since they are focused on expertise. However, I think it is strongly plausible that practice, in some form, is also needed to develop a capacity even to lower levels. The practice might not be as deliberate, extended, or disciplined, but it is very improbable that it would not be necessary at all. I take it that these studies suggest that practice, in general, is a necessary condition for capacity development.

Some recent philosophical work on capacities also supports this view of capacity development.¹³⁹ McGeer (2018) argues that acquired capacities (e.g. speaking a language, constructing philosophical arguments) need practice in order to be sustained and developed. Agents:

do so by repeatedly manifesting some approximation of the dispositional property in question, and then reshaping how they behave — hence, their own intrinsic features¹⁴⁰ — in light of feedback they receive from the environment. (McGeer, 2018, p. 361)

Practical deliberation is also a capacity. If the above discussion is correct, then it means that we will develop this capacity only if we practice it. Practice

¹³⁹ Some not at all recent philosophical work also supports this idea: Aristotle thinks that this is how skills develop; and, also, how virtue develops (see, e.g., NE II.1; Russell, 2014, p. 19).

¹⁴⁰ The intrinsic features are features of the brain, i.e. cortical networks (McGeer, 2018, p. 362).

will be necessary for developing practical deliberation as well. What does it mean to develop the capacity for practical deliberation? I take it that it means that we get better at deliberating, just as we get better at playing the piano when we practise, rather than just watch how others do it. We become better decision-makers: we increase our sensitivity to the relevant inputs; we correctly select the features relevant to our situation more often and with more ease; we weigh up considerations in a better way; we imagine the possible lines of actions and their consequences more easily and with more accuracy; we learn more about what our normative commitments mean and how it is best to apply them.

How do we develop the capacity for practical deliberation? Well, just like with the other capacities, we do that by exercising it; that is, by engaging in processes of deliberation when the occasion presents itself. When we are confronted with making a decision and we try to figure it out on our own we are seizing an opportunity to recognise the reasons that apply to us; to weigh them up against one another; to imagine what would happen if we did this or that; to specify our values. This is how we will be able to better deliberate in the future. However, if we defer recurrently, we stop taking the opportunities to deliberate on our own, because we get other people's testimony to solve our conundrums instead. We replace engaging in processes of practical deliberation ourselves with deference. This way we stop exercising our practical deliberation. And if exercise is necessary for the development of this capacity, then, by deferring, we interfere with its development.

Of what does this interference consist? We simply do not get better at finding out the relevant reasons and the normative and non-normative features relevant to our questions, at imagining possible scenarios and their implications,¹⁴¹ at specifying our values in the right way, and so on. Just like a piano player who stops practising playing the piano does not get better at playing the piano and does not develop this capacity further. This would mean that we would not sharpen up our sensibilities such that we become better deliberators. Instead, we will be *less* likely to know how to deal with new situations where we cannot decide what to do automatically. This can translate into making our capacity for practical deliberation less reliable, less safe, and narrower in

¹⁴¹ Kind (2019) defends the claim that sensory and experiential imagination are themselves skills that can be developed.

application. We would not lose it altogether, and we might even keep some specialised parts of it (e.g. we can still be good at choosing to whom to defer), but we would have the strength of our sensibilities diminished rather than increased. For example, we would not pick out as easily or as often the relevant reasons or salient features of a situation; we would be less likely to be right in the future in our normative calculations; we would not be able to move from deliberating about one thing to a different thing, say from moral to the prudential, because we would not be good at applying our deliberation to diverse matters.

Moreover, I think a case could be made for the claim that not exercising the capacity for practical deliberation can cause it to deteriorate, to become rusty and atrophy. Again, this seems in line with many of our intuitions and practices. It is probably why we make a point of exercising capacities we do not want to lose. Think about the ability to speak some language we learned in school. If we have not had a chance to practice it, we probably are not so good at it anymore. For example, I want to keep my ability to speak and read French and not have it get rusty, so whenever I have some free time, I watch the news in French or I practise my speaking on different apps. This typically works, which gives us at least some justification to believe that the manifestations of a capacity (their occurrence or lack thereof) have an influence on the state of the capacity — that is, regarding its future exercise, i.e. whether it will develop or not, become more or less reliable. This pattern appears to apply to many of our capacities, both cognitive (e.g. solving crossword puzzles; speaking language; memory; chess) and physical (e.g. musical instrument playing; some sports).

In the empirical domain, research on the decay of skills is scarce. Some studies suggest that high or expert-level capacities tend not to worsen despite disuse, or at least that very little retraining will get agents back at the same level (Ericsson et al., 1993, p. 388). However, this seems to hold only for capacities developed through deliberate practice. By contrast, capacities in which expertise and deliberate practice do not come into play that much, for example the capacity to solve crossword puzzles (Moxley et al., 2015), do tend to worsen if unused.¹⁴² The difference between the two situations seems to be the presence

¹⁴² The authors put solving crosswords puzzle into the category of games and contrast it with professional activities, such as sports, chess, music.

of deliberate practice in one but not in the other. This might be accompanied by other factors, such as the distinction between professional work and play, or having the performance of the capacity as a goal in one's life or not. If this is true, then the capacity for practical deliberation would be, in these respects, more like the capacity for solving crossword puzzles rather than the capacity for music. The capacity for practical deliberation, too, is not typically something that is developed with the aim of making it a professional activity, nor is it subject to extensive deliberate practice, in the sense specified above. It does not seem like the kind of thing we do in order to explicitly achieve expertise. This similarity could mean that the capacity for practical deliberation, just like that for crossword puzzles, could worsen if not exercised.

Again, McGeer's work supports this. She argues that acquired abilities need practice to be sustained, as they get rusty otherwise:

We are all familiar with this phenomenon: think of the sports or musical instruments you could play long ago, the algebraic or calculus problems you were able to solve with ease, that second (or possibly native) language you could speak and understand with relative fluency, perhaps even the bicycle you used to ride with no hands and in utter confidence of your ability to stay upright. Alas, no more. (McGeer, 2018, p. 362)

This is explained, she argues, by the fact that these capacities, as dispositional properties, are constituted at least partly by some features of our brains (probably cortical networks) and are fragile. That is, they are maintained only if used; otherwise they decay. Capacities are dynamic, they need to be practised to be sustained and developed, but also to prevent them from getting out of shape. If this is true, there is no reason to think capacity for practical deliberation does not follow this pattern.

So, if we do not exercise weighing and taking into considerations different aspects of a situation or appreciating the reasons that apply, we might be less able to do that well in the future because our capacity would decline. We would not know where to look for relevant information, what counts as relevant and how we go about finding that out; we would not be sure what courses of actions matter and would not be able to imagine how they could go. For example, sometimes recognizing what aspects are morally relevant requires

paying attention to other people and their reactions. Someone who stops attending to others could become less sensitive at registering when and how much another person's interests are salient, at handling other people's feelings, at responding depending on the person's circumstances etc. In short, our sensitivities could get rusty. At best, it would be difficult or very difficult for us to engage in good processes of practical deliberation and would run the risk of them being incomplete and poorly managed. At worst, the capacity for practical deliberation could decay and become less reliable.

The claim that not exercising the capacity for practical deliberation might lead to its deterioration depends, perhaps, on further empirical research. My argument does not stand or fall with it, but it is worth keeping in mind that, if it is true, recurrent moral deference can interfere with practical deliberation in far stronger ways than we have thought. But, for now, *the interference thesis* states that recurrent moral deference interferes with our capacity for practical deliberation, in both its exercise¹⁴³ and its development.¹⁴⁴

This is, then, the argument against recurrent moral deference. Recurrent moral deference interferes with the exercise and the capacity for practical deliberation. The capacity for practical deliberation is both instrumentally and (extrinsically) finally valuable. Given that it is *pro tanto* bad to interfere with something of value, recurrent moral deference is *pro tanto* bad when and insofar as it interferes with the exercise and development of practical deliberation.¹⁴⁵

¹⁴³ It is worth emphasising that, although I mention here the exercise of practical deliberation, this does not mean that interfering with one instance of the exercise of practical deliberation is *pro tanto* bad. My claim is about recurrent deference, so only repeated interference with the general exercise of practical deliberation is problematic.

¹⁴⁴ One might say that deference is itself needed to develop practical deliberation, even as adults. This is because, besides exercising this capacity, we also need to calibrate it, which can involve recurrent deference. For example, we might make a decision about what we should do, but then we also seek to see whether an expert would come to the same conclusion. Now it is unclear why this would call for *deference*, rather than further discussion. We can compare our judgements to other people's judgements without deferring. We can collaborate with others without adopting their beliefs. If this is not accessible to us, but we *need* to defer, then deference would be all-things-considered problematic. But if we defer to the extent that we interfere with practical deliberation, then it will be nevertheless *pro tanto* problematic. I thank Paulina Sliwa for this objection.

¹⁴⁵ One question that might further arise about my view is: is it not the case that my argument overgeneralises to *all* cases of deference? Perhaps it is possible to construct a similar argument with respect to theoretical deliberation. I am not entirely sure such an argument would work. But, as we will see in the next chapter, I am not opposed to my view extending to other kinds of deference. However, answering this question will have to be left for a future project.

Before, I end this section it is worth noting a question that could arise here. I say *recurrent* moral deference is *pro tanto* bad, but how many times is too many times? In other words, when do we have recurrent moral deference rather than a few isolated cases? It is notoriously hard to set thresholds, so giving an exact number would be difficult, if not impossible. But an exact solution might not be necessary. If we believe that practice is indeed needed to be able to exercise and develop one's practical deliberation, then we can also believe that there will be a point where there is such a thing as too much deference. So, I will only say that one has deferred too many times when one has deferred enough to compromise one's capacity for practical deliberation, however many times that may be. When our sensibilities are so affected that they are less good at doing their job anymore, then we have deferred too much.¹⁴⁶

2. Objections

To sum up, so far, I offered a systematic explanation for the problematic nature of moral deference. I said that what is *pro tanto* bad about recurrent moral deference *per se* is that it interferes with the exercise and the development of a capacity that is valuable, i.e. the capacity for practical deliberation. Repetitively circumventing the exercise and development of our practical deliberation will prevent us from becoming good deliberators and acting in accordance with the existing reasons, as well as from shaping our normative commitments and constructing our view of the normative world and our conception of how we should live. Now I consider some objections.

2.1. The Overgeneralisation Worry

Given how I have just described my account, one might wonder whether it implies that any activity that is done *instead* of deliberating, at any point, is *pro tanto* bad? Another way to put this worry is: since practical deliberation is so

¹⁴⁶ This also provides an answer to the problem of mixed moral deference, i.e. cases where the moral aspects of a problem are inseparable (in practice or in principle) from its non-moral aspects. According to my account, interference with practical deliberation is doing the work, regardless of whether we can strictly separate the moral from the non-moral in every case.

valuable, we may have reason to constantly promote it, and make sure we exercise and develop it. For example, I could be doing some deliberation right now instead of writing; or I could be deliberating instead of doing other daily activities, like reading, jogging, sleeping etc.

My answer is no. First of all, there are plenty of valuable things out there and we do not focus exclusively on promoting them all the time. Pleasure is valuable, but we do not seek it constantly. Instead, many times we actively and voluntarily put it aside and do something that does not bring us pleasure. And we do not think that this is *pro tanto* bad. It does not follow from something being valuable that it needs to be constantly promoted and that it should replace everything else.¹⁴⁷

Secondly, insofar as practical deliberation is a way of recognising and responding to reasons and particular features, it will be connected to specific situations where it is an *appropriate* thing to do. Simply reading or jogging is *not* connected to situations where practical deliberation is an appropriate thing to do. However, maybe they could be. Imagine Jack. Whenever he has to make a decision — and as such it would be appropriate for him to use practical deliberation — he goes to sleep instead and refuses to make the decision. He replaces practical deliberation with sleeping. Is sleeping *pro tanto* bad then? I am going to bite the bullet and say yes. Insofar as, because, and in virtue of, supplanting practical deliberation where it would have been appropriate to use it, sleeping is *pro tanto* bad *in that respect* and *to that extent*. This is just the other side of the coin of believing that practical deliberation is valuable in situations where it is called for.

Now one might have a further worry here: it turns out that on my view there is nothing distinctively problematic about (moral) deference, since even sleeping can interfere with practical deliberation in certain special cases. First of all, if this is a problem, it does not only challenge my view. Recall that all other pessimist views are based on the importance and value of some thing: moral understanding, virtue, appreciative normative knowledge, authenticity. These

¹⁴⁷ There are many attitudes and responses we could have towards values besides promoting them, such as (but not limited to) preferring, respecting, cherishing, or caring for them (Rabinowicz and Ronnow-Rasmussen, 2000, p. 46). That interfering with practical deliberation is *pro tanto* bad need not amount to saying that it should always be promoted or maximised. My preferred way of thinking about it is that it is *pro tanto* good *to respect* this value (by exercising and developing it rather than let it decline).

things may well be interfered with by, e.g., sleeping in the same way. Secondly, some of the pessimists have already admitted that deference is not the only way in which the values that their accounts are based on can be undermined. Take Fletcher (2016), whom I will discuss in chapter VI. His view is that moral deference is problematic because it does not generate the appropriate moral sentiments in the agent who is deferring:

I hope it is already clear that I do not claim that moral deference is the *only* way in which one can end up with a problematic disconnect between one's moral beliefs and one's moral sentiments. The same problem occurs for agents who, themselves, form moral beliefs but whose conflicting moral sentiments are recalcitrant, such as the person who convinces themselves that meat-eating is wrong but cannot blame those who do so. (Fletcher, 2016, endnote 40)

Mogensen (2015, p. 264) makes even a stronger point: “What we have to keep in mind is that pessimism is as much—if not more so—a view about the importance of figuring things out for ourselves.” This suggests that the distinctive problem is not that people defer, but that they do not figure things out on their own.

Thirdly, there does not seem to be a particular reason why we should think that (moral) deference is distinctively problematic. Our intuitions do not necessarily seem to point at that. They suggest that there is something problematic with (moral) deference, but not that there is something necessarily *uniquely* problematic with it. But I do not see this as a problem. Why would it be? Our task was to figure out whether our intuitions are warranted and explain why moral deference is problematic — and this is still an interesting issue. The fact that other phenomena may be problematic in the same way does not take away from that.

2.2. Another Overgeneralisation Worry

We often defer to our doctors about what medicine to take; perhaps to our friends about how to cook certain meals; to different trustworthy people on questions such as whether to cook or to do laundry, where to shop, what cereal

to buy, how to work our new washing machine etc. There are plenty of mundane practical issues that we defer about and it does not seem to be problematic. My view, however, might imply the opposite: this is deference about practical matters and so it interferes with practical deliberation. To that extent, it should be deemed *pro tanto* bad.

My response is the following: recurrent deference about trivial, everyday matters is not *pro tanto* bad because, although practical in nature, it does not interfere with the capacity for practical deliberation. On the one hand, we have the everyday matters that we could not have deliberated our way to: what kind of medicine we should take when we are ill; how to work this new electronic device I got; how to fix my broken pipe. In such cases, no matter how hard I tried, no matter how much I deliberated, I simply lacked the necessary information.

On the other hand, we have the matters that we could figure out on our own but we consider that we have better things to do: When we defer about the order in which we should do our chores, what to have for dinner, whether to buy a new vacuum cleaner etc., we do it mostly for convenience. We want quick answers; we cannot be bothered thinking about them properly because they are mainly trivial, unimportant matters. We defer because it makes our lives easier. But precisely because such issues are trivial we could, at any point, deliberate on our own and find a solution. Moreover, it might not even count as deference a lot of the time. I do not defer to my partner about what to have for dinner because he is an expert in dinner-menu-choosing. Rather, I do not want to think about it now, and maybe I do not care too much. So, I ‘defer’, in the sense that I leave it up to him. Deference involves taking someone as better equipped to deal with a certain problem. In this case, we are both equipped to decide what to have for dinner. I am just too lazy to use my equipment right now. Thus I would say deference in these instances does not interfere with practical deliberation. My account deems these cases unproblematic because one of the two elements of my account is not seriously fulfilled. Although they fall under the jurisdiction of practical deliberation, such questions do not interfere with it.

2.3. Does Practical Deliberation Matter for Moral *Beliefs*?

While the focus of my argument is on practical deliberation, one can defer both by acting and by forming beliefs on the basis of someone else's moral testimony. If someone recurrently defers about moral judgements rather than actions, is their capacity for *practical* deliberation interfered with? Does not practical deliberation result in intentions or actions, whereas beliefs — including moral beliefs — result from theoretical deliberation?

Firstly, practical deliberation is usually considered practical in virtue of both its upshot, i.e. it ends in an intention or action, and its subject, i.e. its content, is practical, concerned with what to do. So, a process of deliberation can count as practical in an important sense even if it does not produce an intention or an action, but it aims at answering the question of what to do. And our moral beliefs tend to be about what to do, e.g. from a first-person perspective ('I ought to ϕ ') or a more general perspective ('One ought to ϕ '). Thus, when one defers about a moral judgement that would otherwise have been formed through deliberation, *practical* deliberation is interfered with. Moreover, even if some moral judgements are not explicitly about what to do, e.g. 'The Nazi regime was immoral', 'Racism is unjust', it is fairly plausible that they are still action-guiding in some sense or that they suggest certain actions or attitudes, e.g. that we ought to condemn the Nazis, that we ought not be racist.

Secondly, in general, moral beliefs are likely to be at least indirectly connected to action, even though not immediately so.¹⁴⁸ For example, think about forming the belief that capital punishment is morally impermissible. For most of us, this belief will not translate into any immediate action. However, even if we are in no position to form any intentions or do anything about it, the moral belief that we form will relate to other judgements which will have direct practical relevance and will be transformed into some intentions or actions. That

¹⁴⁸ This is loosely based on a point that Nickel (2001, pp. 260-261) makes in a different context. He argues that our moral beliefs should be accompanied by moral understanding because, even though they themselves might not be immediately translated into action, they will likely be connected to other moral beliefs that will be put into action. The thought is that if we form a belief by deference, which itself is not relevant to action, and that belief will be connected to another belief that is immediately relevant to action, we will not have moral understanding of the latter and thus our action will lack moral worth. I do not subscribe to the details of his point, but I do think that it is important that he draws attention to the relation that exists between different moral beliefs.

is, it will be connected to moral statements concerning the value of a human life irrespective of what humans do, or how we should punish people who do bad things. It might also influence our behaviour and interactions with people who have been convicted of different offences, or it might make us support or retract support from politicians who have similar or dissimilar views to ours. In other words, besides being practical in content, the moral judgement will also be at least indirectly practical in its upshot, as it will serve as an input into a ‘purer’ practical deliberation process, i.e. one which will result in an intention. Moral judgements will rarely, if ever, be completely practically irrelevant. Given all of this, I want to hold that recurrent moral deference, even if it is about moral judgements, still interferes with *practical* deliberation.¹⁴⁹

2.4. Is Recurrent Deference Truly Incompatible with Deliberation?

Someone might think it is unclear whether it is plausible that recurrent moral deference interferes with the exercise and development of our capacity for practical deliberation. This is because: 1) moral deference might be recommended by practical deliberation; and 2) one can defer recurrently and thus fail to exercise and improve one’s capacity for practical deliberation in one domain, e.g. the moral, but exercise it and improve it in another domain, e.g. the prudential. Let me take them in turn.

2.4.1. Moral Deference Might Be Recommended by Practical Deliberation

Regarding objection 1), I agree that practical deliberation might sometimes recommend moral deference. However, first, I think this is limited to urgent and high-stakes cases. And, second, that moral deference is recommended, i.e. all-things-considered unproblematic, does not eliminate the

¹⁴⁹ A different way to develop the thought that moral beliefs are connected to action is via theories of rationality or metaethical theories. For example, one might say that if one believes one ought to ϕ , there is at least some rational pressure for one to intend to ϕ (e.g. Broome, 2013). Or that if one judges that one ought to ϕ , one will be motivated, at least to some extent, to ϕ (e.g. Smith, 1994).

possibility of it being *pro tanto* problematic. For instance, take Enoch's (2014) example, where he is uncertain about whether to vote to fund a war which will impact many lives and is thinking of deferring to his friend who seems like an expert on this kind of thing. Here, Enoch might deliberate about his and his reliable friend's epistemic states, reflect on the consequences of each line of action that he could take, and then decide that it is better (or morally required) to defer. In situations like this, where one needs to act, but is uncertain, temporarily impaired in some way (biased, angry, drunk etc.) or pressed by time, practical deliberation can indeed recommend deference. Enoch (2014, p. 258) says moral deference might even be morally required in such a case.

But this is not an issue for my account, because I can agree with Enoch that one should all-things-considered defer here. My claim is *pro tanto*, so I accept that moral deference might be the wise solution or that it might outweigh the value of non-deference sometimes. Plus, I also think that only recurrent moral deference is always *pro tanto* bad. When it comes to other scenarios, where these features of urgency are not present, then it is not clear that practical deliberation would indeed recommend deference. Instead, perhaps we should be agnostics and suspend judgement instead. This is especially so if we think that an unclosed process of inquiry entails such an attitude. As long as a question is open (i.e. unresolved) and the process of inquiry is ongoing (albeit not necessarily active), one should plausibly continue to suspend judgement (Friedman, 2017). If there was a reason to defer rather than to suspend judgement in such situations, then the burden is on those who claim that to show why.

Still, it might be that this objection hints at something deeper: moral deference is not strictly incompatible with the exercise and development of the capacity for practical deliberation. For one, we can defer but still deliberate, in some way or another. For example, someone can still exercise practical deliberation when one deliberates about whether and to whom to defer. Take fictional Enoch, who comes to defer following a process of deliberation, in which he assessed his friend's track record as well as his own, and he weighed the consequences of making a decision alone. Strictly speaking, I agree that deference is compatible with practical deliberation to a certain extent. We can certainly deliberate about whether or not we should defer, and then end up deferring. However, I hold that, given how limited the scope of practical

deliberation is in cases where we just defer about whether to defer and to whom, the capacity will still be interfered with if one recurrently defers.

For example, take unlucky Alan, who always ends up in urgent situations of moral uncertainty, and correctly deliberates and concludes that he ought to defer. He is good at deliberating about what to do when he does not know what to do, and chooses his advisers wisely, but he cannot deliberate about the things that his advisers deliberate about. I think it is plausible that we would not say that his capacity for practical deliberation is very good. Engaging in successful meta-deliberation — deliberation about when to deliberate or what to do when you cannot deliberate — does not amount to good general practical deliberation. Someone like Alan would not be able to navigate his practical world on his own; to figure out what to do by himself; not to mention that he would not be able to set his own normative goals and to specify his values. But this is precisely the business of practical deliberation and probably the reason why good meta-deliberation is not the paradigmatic example of practical deliberation. Meta-deliberation can, of course, be very important. It is good to recognise our own limitations, to be able to figure out when we should defer, as well as how to choose our experts. However, our capacity for practical deliberation is not valuable merely because it can help us do that.¹⁵⁰ It would be a rather limited capacity if it only or mainly helped us to figure out whether and to whom we should defer.

Another case could also be used to make the point that practical deliberation is not incompatible with deference. We can imagine individuals whose capacity for practical deliberation is fairly good; or perhaps they are even perfect deliberators. They nevertheless choose to defer rather than do their own deliberation. However, recurrent deference will not interfere with the development of practical deliberation of these agents, as their capacity is already developed. Presumably, moral deference will still trigger negative intuitions. But can my account explain why? I have two answers to this. First, if the deterioration claim is correct, then my account has the resources to explain why

¹⁵⁰ Another way to reply is to say that the problematic character of moral deference comes in degrees. It seems worse if one curtails practical deliberation completely and just defers about the answer one is seeking than if one defers only about what considerations are morally relevant. So, it is true that moral deference might not be strictly incompatible with practical deliberation, but the capacity will still be more or less hindered, depending on the kind of deference that happens.

such individuals recurrently deferring might still seem problematic. Specifically, if it is true that not exercising one's practical deliberation can lead to its atrophy, then this is why it is *pro tanto* bad to recurrently defer, regardless of how developed one's capacity already is. The second response is that such cases might actually not be problematic. Or, at least, we could not say whether they would be problematic *per se*. If practical deliberation is not interfered with, then recurrent deference is not *pro tanto* bad. However, it may seem like it is, especially by looking at individual cases, because a question arises: if someone who has good practical deliberation and is able to figure things out for themselves, why would they defer? This might point to some character flaw existing previously to the deference. So, if their deference seems fishy, it might be because it points to something that is wrong with the person already (depending on why they defer). But this does not point at anything that is wrong with deference *per se*.

And, finally, there is another possible situation that could pose some problems to my account. We could have people who do their deliberation, reach a conclusion, but then for one reason or another they defer. This way they exercise and practise their practical deliberation, but nevertheless end up deferring. What does my position say about them? Let us imagine how the situation would unfold: John is deciding whether it is morally permissible to lie in a certain situation. He deliberates on his own and reaches the conclusion that it is not morally permissible to lie in that situation. Then he decides to ask an expert. The expert could either agree or disagree with him. If the expert agrees, then there is no need for deference, as John can sustain that belief based on the reasons that ground it. If the expert does not agree, then John might think that, since the expert knows better, he should defer. We might all believe that it is indeed all-things-considered better than he defers, given that an expert is, by definition, better placed to know the answer. This happens numerous times regarding numerous issues. Is this recurrent deference nevertheless *pro tanto* bad? Yes. It is clear that John's capacity for practical deliberation is still in need for development, so deference, instead of more deliberation and consultation with others, would be *pro tanto* bad.¹⁵¹

¹⁵¹ This gets us to a related problem: the issue of *active deference*, i.e. deference in action and not in belief. For example, I give money to a homeless person because you, an expert, told me to do

2.4.2. Deferring in One Domain, Deliberating in Others

Now I turn to objection 2), which says that it is plausible that one can defer recurrently, and thus fail to exercise and improve one's capacity for practical deliberation, regarding a certain type of matters, e.g. the moral, but simultaneously exercise it and improve it with respect to other type of matters, e.g. the prudential. Or it could even be that we actively want this kind of division of labour when it comes to practical deliberation. For example, we can organise so that each of us becomes an expert regarding a particular domain or issue, such as the value of animals, or matters of war. By designing this division of labour, we can be more reliable together. This then suggests that recurrent moral deference might be compatible with a well-functioning capacity for practical deliberation — or even that such deference is desirable.

However, I do not think this would be a good result. Firstly, relying on a division of labour when it comes to practical deliberation would mean that each of us would be specialised and reliable in a certain area, regarding a specific problem, or some procedural aspect of practical deliberation. But, although reliable as a group, our individual reliability would be narrow. That is, this kind of division of labour would restrict the range in which each of us can reach practical conclusions. Each of us would then be utterly unresponsive to a range of different kinds of reasons that exist. So, my capacity for practical deliberation would not be practically safe or robustly reliable, as I could not navigate my practical life without the help of other, differently specialised, individuals. Thus, even if recurrent deference were to occur only in a certain domain, it would still ensue in a kind of interference with our capacity for practical deliberation, namely an interference in its reliability or safety. The capacity could still develop (in a very specific way), but as a general capacity it could not be very broadly applied and as such it would not be reliable or robustly safe.

so, without forming a judgement that this is the best thing to do. Active deference, by definition, does not preclude practical deliberation. So does recurrent active deference count as problematic on my view? It all depends on how such deference comes about. If I have not deliberated at all, then yes, it is problematic because it interferes with the exercise and development of practical deliberation. If I have deliberated, then we go back to what I say in the paragraph to which this footnote belongs, and my response will be the same.

Developing the capacity in such a narrow way has certain implications given the capacity's instrumental value discussed above. For one, we could not figure out what to do in areas other than our own specialisation if the experts would not be around. We simply would not be able to answer all the practical questions that we would be confronted with. Secondly, our future selves' activities would also be endangered, for the same reason. Moreover, there are implications given the (extrinsic) final value of practical deliberation. If we localized our expertise, we would not be able, on our own, to specify our own values, define our commitments, qualify our ends, or build a conception of how to live. We could not lead that kind of life which seems valuable. With respect to this feature of practical deliberation, i.e. being the procedural skeleton of this sort of life, it would always be *pro tanto* bad to recurrently defer, if recurrent deference interferes with the good, safe and reliable, functioning of the capacity.

Although it is true that our limitations might sometimes give us reason not to specialise in all aspects of morality — perhaps it is simply impossible for us all to be experts when it comes to ethical issues related to animals, sexism, war, mental health, environment, child-rearing, and so on — it does not mean it is not *pro tanto* bad to interfere with the capacity that can get us close to that, namely with our practical deliberation.

3. Concluding remarks

In this chapter I have argued for a novel version of moral deference pessimism. I hold that recurrent moral deference is *pro tanto* bad because it interferes with the exercise and the development of our capacity for practical deliberation. This capacity seems to be both instrumentally and (extrinsically) finally valuable and by hindering it deference becomes a problematic phenomenon. In what follows I want to tie up a few loose ends that my argument might have, and conclude by mentioning some of the advantages it could have over the rival pessimist views.

To return to an issue I raised in chapter II, it is worth mentioning how my account handles individual instances of moral deference. As I have argued there, my view is that such cases are not fruitful to investigate as they cannot show us anything about moral deference *per se*. I have also said that since no

ideal or value requires perfect compliance individual cases are not *pro tanto* bad either. This is upheld by my practical deliberation view. One-off interference with practical deliberation will not have sufficient negative impact on this capacity; which is exactly what I say about all other explanations of the problematic character of moral deference. At best, we can say that individual cases of moral deference are problematic when, and insofar as, they are part of some pattern of cases that is *pro tanto* bad.

I also have to say a few words about agnosticism and deference. Hopkins (2007, p. 613) and Mogensen (2015, p. 270) mention the fact that agnosticism seems a more appropriate response than deference. That is, in a situation where we do not know what to do on our own, it is said, the intuition is that it is better to remain agnostic and suspend judgement, rather than to defer. Mogensen (2015, p. 270; p. 273) even faults the other existing pessimist views on the basis that they do not render agnosticism as preferable to deference. So how does my view see agnosticism?

My view, of course, centres on practical deliberation. So, if there is a choice between remaining agnostic or deferring, that means that the answer has not been reached and more deliberation is appropriate (i.e. in order to get that answer). Then, agnosticism is preferable to deference because it is more likely to allow for that further suitable deliberation to happen. This is because deference would settle the question and plausibly undercut the possible deliberation that could have fittingly happened down the line.¹⁵² Agnosticism, by contrast, leaves the door open for additional deliberation. The thought is that one might come across new information that one needed to undertake that deliberation that would have given them the answer. This might then enable one to go back to the question and deliberate more such that they get their answers on their own. This way, they would have the opportunity to exercise and develop their capacity. But if they choose deference, the question is put to bed, as they immediately get the answer they wanted, undercutting the deliberation which would have been appropriate. This makes it unlikely that they will go back to

¹⁵² Note that this is similar to Callahan's (2018, pp. 455-456) claim that deference dis-incentivises the acquisition of understanding. However, we use it for different things: she wants to say that this is the reason why moral deference is in tension with obtaining understanding.

deliberating, and if this is done recurrently it interferes with the exercise and development of the capacity.¹⁵³

Finally, to end, I want to mention a few advantages that my view may have over the existing pessimist accounts. Firstly, as I have said in the beginning, there are potential structural parallels between my argument and some of its rivals. Perhaps the moral understanding strategy could be reconstructed to say that recurrent deference interferes with the capacity for moral understanding. With that in mind, I want to highlight one of the most important features of my account, which can set it apart from other views. The practical deliberation view naturally supports a version of the asymmetry thesis, the claim that there is a difference between deferring about certain matters but not others, such that some types of deference are problematic, but others are not. It does so in a way that is better than what the other accounts can offer. As I will show in the next chapter, the existing literature has misconceived the asymmetry thesis and it is mistaken in delineating it according to domains, e.g. moral versus non-moral domains or normative versus non-normative domains; this does not track the problematic versus unproblematic deference distinction. I take my argument to provide a better criterion for drawing the asymmetry: interference with practical deliberation. As such, my view is able to provide a more unified explanation for the seemingly problematic character of other types of deference, e.g. the aesthetic and the prudential, and it can also give the right predictions in the various normative and non-normative cases, e.g. legal, etiquette. I take this to be a significant step in the debate about deference.

Secondly, my argument offers an alternative to rejecting moral deference which better recognizes our human limitations and does not ask too much from the agent. For example, it is less intellectualized than views which explain the seemingly problematic nature of moral deference in terms of a deficiency of

¹⁵³ Mogensen (2015, p. 271) poses this objection to Hills (2009) and imagines a reply from her that is similar to mine (except that it concerns moral understanding). He rebuts it by saying we can imagine cases where the possibility of gaining moral understanding later regarding a particular question is not connected to whether or not the person defers now. That is, imagine we know they will gain that understanding, even if they defer now. So it does not matter whether one defers or remains agnostic now. Can this response apply to what I have said? Perhaps. I think it is not that problematic if I bite the bullet here and say that if the person knows they will have the chance to exercise their deliberation and reach the conclusion on their own later, in this pattern of cases, then they are indifferent between agnosticism and deference. But is the intuition that agnosticism is preferable to deference still standing, given that we have knowledge that deference will not sustain the beliefs/actions in question for too long?

moral understanding (Hills, 2009; Hopkins, 2007; Nickel, 2001). As we have seen, the best developed version of this view takes such understanding as *understanding-why* (contrasted with knowing-why), which is a primarily explicit kind of understanding here, involving abilities such as being able to follow explanations of why some moral judgement obtains as well as to articulate such explanations (Hills, 2009, pp. 102-103). Without this understanding, our actions have no, or little, moral worth. By contrast, my account acknowledges that we will occasionally need to defer; that it is likely that we will not be perfect deliberators; and that we might not always be able to articulate the processes of deliberation in which we engage. My view is not about how we fail to reach an ideal, but about promoting that we do the best that we can and showing why that is important.

Thirdly, my version of pessimism recognizes more explicitly the legitimacy of the optimists' intuitions and claims. I take this to be important because we need to accept that testimony-based beliefs and actions are inevitable, as well as useful, even when it comes to moral matters. If we want to tell a plausible story about moral deference, we must consider a greater variety of intuitions and acknowledge what the optimists get right. I believe my argument can do just this, while still keeping the rivalry between optimism and pessimism and vindicating our negative intuitions.

And, finally (this time for real), I also want to emphasise that my account does not imply that there is no space for collaboration between people and that we should never talk things through with others or ask for help. On the contrary, I think my position can recognize the importance of joint deliberation and the contribution that other people can make to our own normative development. We can learn so much from others. And what we learn from others, the arguments they have — it can all be used to shape what we ourselves think about those issues. It provides us with new information, it helps correct mistakes, and it expands our world view. But this is not incompatible with holding that individual practical deliberation is also important and valuable. After all, my claim is a *pro tanto* one. The pessimist does not need to knock over the glass just because it is half empty.

Chapter V: The Asymmetry of Deference

In the previous chapter I presented my own version of moral deference pessimism. I argued that recurrent moral deference is *pro tanto* bad because it interferes with the exercise and the capacity for practical deliberation. But the keen eye will have noticed by now that practical deliberation does not help us solve just our moral questions and conundrums. Rather, the jurisdiction of practical deliberation extends to, well, all that is practical. Thus, my argument will also extend beyond moral deference to practical deference in general. This is to say that if some pattern of recurrent deference is problematic, it is so insofar as, to the extent that, and because it interferes with the exercise and the development of our practical deliberation. This is what makes it *pro tanto* bad. I will now argue that this implication of my account will help us make sense of, and better outline, the corollary claim that the typical moral deference pessimist holds, namely *the asymmetry thesis*.

The asymmetry thesis is the claim that some instances of deference are (non-epistemically) problematic *per se* while others are not.¹⁵⁴ Philosophers have mainly suggested that we can distinguish between problematic and unproblematic deference by looking to what domains¹⁵⁵ of discourse the matters we defer about belong (e.g. Hopkins, 2007; Lord, 2018). More specifically, the existing ways of accounting for the asymmetry propose mapping the problematic-unproblematic distinction onto the moral-non-moral or, alternatively, normative-non-normative distinction. Generally, we find moral or normative deference problematic and non-moral or non-normative deference unproblematic (e.g. Hills, 2009, Howell, 2014).

This chapter investigates whether there is in fact such an asymmetry and, if there is, what shape it takes. In what follows, I put forward an argument that has two parts: firstly, I argue that the domain-based asymmetry thesis is wrong and, secondly, I propose a different basis for a new version of the asymmetry thesis. For the first part, I will examine various examples of deference to show

¹⁵⁴ As always, throughout this chapter too, I have in mind deference *per se* (even if I do not explicitly say that) unless stated otherwise.

¹⁵⁵ I use the term ‘domain’ for brevity, to illustrate the distinctions the defenders of the asymmetry thesis make. However, I acknowledge it is an unclear and unhelpful term, as I will show later.

that a domain-based delineation makes the wrong predictions. The asymmetry turns out to be a lot less tidy, as the distinction between problematic and unproblematic deference cuts across domains. For the second part, I propose that the practical deliberation view developed in the previous chapter can provide a basis for the asymmetry thesis in a way that is in broad accordance with our intuitions.

In turn, this will make my practical deliberation view the most comprehensive view on deference that currently exists on the market. It will offer a unified account of the non-epistemically problematic character of deference in general and it will not be limited to any particular domain(s). As mentioned in the previous chapter, the other pessimist accounts could, in principle, restructure their views and answer the objections others and myself have posed to them. Given that, the virtue of being the argument that manages to best explain the issues of deference in the broadest way will make my practical deliberation view at least *pro tanto* preferable to its rivals.

1. The Domain-Based Asymmetry Thesis

To remind ourselves what the asymmetry thesis is, think about how we often defer about what the time is, where events take place, information about current states of affairs, weather predictions, or history, chemistry, physics, and medicine-related facts. Such deference strikes many as indispensable and unproblematic. Yet not all deference is like this. The intuitions shift when it comes to some other kinds of deference. It seems suspicious or problematic in some way to defer about things like whether capital punishment is morally wrong, what career paths we should take, or about the beauty of a painting.

This sort of intuitions has been the basis for a widely made assumption in the moral deference literature, namely that there is an asymmetry between different types of deference. *The asymmetry thesis* (Groll and Decker, 2014, p. 54) claims that some types of deference are (non-epistemically) problematic *per se* while other are not. The asymmetry thesis, however, has not been so far expressly addressed and analysed in a systematic way. One reason for this might be that the alleged asymmetry in our intuitions is used only as data to motivate accounts which argue that something is bad about moral (or aesthetic) deference.

Those who hold this would then be more preoccupied with defending moral (or aesthetic) deference pessimism rather than providing an analysis of the asymmetry thesis.¹⁵⁶ I think this might be a perfectly good sociological explanation for the lack of interest in the asymmetry thesis specifically. However, the asymmetry thesis in itself is important, as well as relevant, for the debate. If we have non-epistemic reasons not to defer about certain kinds of questions, e.g. moral, should we not investigate whether these reasons apply to other kinds of questions? Is it true that only some instances of deference are problematic, whereas others are not? How do we systematically distinguish between them? For one, these are interesting questions in themselves. Two, if we think our views about moral deference recommend certain behaviours and ways of forming beliefs, then it seems significant to investigate whether other kinds of deference do that too. So in what follows I will do just this. I start with my own reconstruction of the issue of asymmetry, with the aim of clarifying the landscape.

The asymmetry thesis, as we will see shortly, is typically drawn according to domains of discourse. That is, deference in some domains (or about matters that pertain to a particular domain) is problematic¹⁵⁷ whereas deference in other domains is not. As I understand the literature, the asymmetry thesis comes in two versions according to what the contrasting domains of discourse are taken to be. *The narrow version* states that moral deference is problematic, but non-moral deference is not. This is how the asymmetry thesis was first formulated in Hopkins (2007, p. 613). He characterizes moral deference optimists as those who see “no difference in kind between moral and non-moral matters when it comes to taking testimony”.¹⁵⁸ The moral deference pessimists, however, do think there is a difference in kind. Others seem to subscribe to this version of the asymmetry thesis as well (Hills, 2009, p. 95; McGrath, 2009, pp. 322-323).

The broad version of the asymmetry thesis states that normative deference is problematic, but non-normative deference is not. Howell (2014) seems to be a

¹⁵⁶ I thank an anonymous referee who reviewed a paper version of this chapter for raising this point.

¹⁵⁷ Where ‘problematic’ is shorthand for ‘at least *pro tanto* bad’, given the definition of deference pessimism which I have constructed in chapter II.

¹⁵⁸ Recall from chapter III that there is at least one exception to this: Wiland (2015, p. 73) supports an optimist version of the asymmetry thesis. He thinks that at least in some cases it is *better* to defer about moral rather than non-moral matters. He specifically discusses deferring to victims of epistemic injustice.

supporter of this view, as he mentions that perhaps deference in any normative domain is suspicious and that his own virtue-based argument might apply to some extent to aesthetic deference too. Similarly, Lord (2018, p. 74) discusses the sub-optimality of both moral and aesthetic deference and holds that there is “an asymmetry between the merits of deferring on moral and aesthetic matters and the merits of deferring on all other matters”. He sometimes explicitly refers to the asymmetry as being between normative and non-normative deference. Fletcher (2016, pp. 57-58) also subscribes to the broad asymmetry thesis but admits that not all types of normative deference might be problematic.

Of course, neither of these versions of the asymmetry thesis needs to uphold only a difference *in kind* between problematic and unproblematic deference. There could be an asymmetry such that some deference — moral or normative — is *more* problematic than other — non-moral or non-normative. That is, the asymmetry thesis could point to a difference in degree. Mogensen (2015, p. 265), for example, denies the difference in kind idea.¹⁵⁹ Also, McGrath (2009, p. 323) can be interpreted this way, as well as Howell (2014, p. 392), and Fletcher (2016, p. 45), since they all use comparative language when they talk about the asymmetry. It is, however, difficult to clearly tell whether all the philosophers I have discussed prefer the difference in kind or the difference in degree view, as they are not usually explicit about it. So although most pessimists seem to subscribe to some version or another of the asymmetry thesis, it is not always clear which kind exactly a pessimist intends to endorse. Moreover, there is also no apparent consensus on what form it should take. I will argue that both the narrow and the broad version fail to do justice to our intuitions, and they cover both too much and not enough at the same time. That is, they end up predicting that cases which seem intuitively unproblematic are problematic and vice versa. Although I think the broad version is more plausible than the narrow one, I will show that neither can provide a principled line of drawing the asymmetry.

¹⁵⁹ Although he refers specifically only to moral and aesthetic deference, and he does not explicitly subscribe to the difference in degree view.

2. Against the Domain-Based Asymmetry Thesis

Take the narrow version of the asymmetry thesis first: moral versus non-moral deference. The problem here is quite simple. If we take the asymmetry thesis to *literally* mean that there is a difference in kind, as Hopkins first formulated it, between moral and non-moral testimony such that the former is problematic while the latter is not, then we would fail to explain why we have intuitions about certain cases that cut across this distinction. There are instances of deference which are non-moral, but are intuitively problematic, and instances which are moral, but are not intuitively problematic. Take some examples from the literature to illustrate this:

NEFERTITI: Hanna just returned from a trip to Berlin during which she saw *Nefertiti's Bust*. Hanna's sister Clara asks her about the museums. Hanna tells her that *Nefertiti's Bust* was especially beautiful. Clara comes to believe that *Nefertiti's Bust* is beautiful solely on the basis of Hanna's word. (Lord, 2018, p. 72)

TRIP: Anna is a journalist who is preparing to go on a reporting trip to a dangerous and conflict-ridden area. She has to tell her family that she will be away but she really doesn't know what to tell them. If she tells them where exactly she's going and why, they will be extremely worried. On the other hand, she worries that by evading the questions she would be lying. She goes back and forth but cannot decide what the right thing to do is and eventually decides to ask a friend whose judgement she trusts. (Sliwa, 2012, p. 178)¹⁶⁰

Intuitively, at least some (and maybe most) of us find the first instance of deference problematic even though it is not about moral matters, and the second one unproblematic although it is about moral matters. There is something fishy about Clara coming to believe that Nefertiti's bust is beautiful just because Hanna told her. But Anna's case seems to not trigger a negative response (Sliwa, 2012,

¹⁶⁰ In case one does not agree that this is an example of moral deference, here is what Sliwa (2012, p. 178) says about it: "Anna might know that her family will be terribly upset if she keeps quiet about the nature of her trip and they will accuse her of having lied should they find out. Nevertheless, she might be unsure whether they would be justified in their accusations. That's because she is unsure about whether what she contemplates doing really does amount to a lie."

p. 178).¹⁶¹ And we can imagine other kinds of non-moral deference cases that are problematic, like prudential deference (Fletcher, 2016), and other cases of moral deference that are unproblematic (Sliwa, 2012). It seems quite clear that the moral-non-moral distinction does not track the problematic-unproblematic deference distinction here.

It seems so clear, in fact, that it is difficult to believe that any of the existing moral deference pessimists would actually hold the narrow version of the asymmetry thesis. Their other works also suggest otherwise. For example, both Hopkins (2011) and Hills (2017b) think that aesthetic deference is problematic too. So perhaps they used the asymmetry thesis just to prove the point about a specific type of deference, namely moral deference. Nevertheless, it was worth explicitly addressing why we should not support the narrow version of the asymmetry thesis anyway: it is too coarse-grained and it does not get many cases right.

What about the broad version of the asymmetry thesis then? The distinction between problematic and unproblematic deference is mapped onto the normative-non-normative deference distinction here. This is a more promising way of conceiving of the asymmetry, as it can at least account for problematic instances of aesthetic and prudential deference. However, I still think it is not successful in aligning with our intuitions. In certain normative cases, deference might actually appear unproblematic. And some non-normative deference can seem problematic too. But before we look at some examples, I need to mention what is meant by ‘normative’ here.

There is a general sense in which any ought claim is normative, as it will contain some kind of a recommendation or requirement. However, a further distinction regarding normativity can be introduced, between what has been called *robust* (or substantive or authoritative) *normativity* and *formal normativity*. Substantive normativity is typically associated with the ought’s of morality (but also those of epistemology or rationality). Formal normativity is typically connected to the ought’s of etiquette, club rules, games, and the law (Finlay, 2019,

¹⁶¹ I am not sure whether other pessimists would concede that this case is intuitively unproblematic. It is plausible to me that it is indeed intuitively unproblematic. But whatever the intuition, pessimists might still say that this is a case where, although all-things-considered good to defer, it is nevertheless *pro tanto* bad to do it. However, I hope to have been convincing in showing, in chapter II, that individual cases may not be even *pro tanto* bad.

p. 205). The distinction can roughly (but sufficiently for present purposes) be understood as follows:

Consider the rules of chess or standards of fashion. We might say that both are “norms” in the following sense: they are standards that can be used to assess whether something (e.g., an action, a style of dress) accords with it. (...) Contrast this thin sense of normativity with a thicker one, which many take to be at the heart of ethics, as well as epistemology. When an agent does something she ethically ought not to do, all things considered, it seems that she has done something more criticizable and mistaken than when she fails to conform to merely formal norms. We invoke this thicker notion of normativity when we ask not just how an agent’s actions stand in relation to a given set of norms she just happens to care about, but rather what she should really do, all things considered. (Plunkett and Shapiro, 2017, p. 48)

So the question is: how does the literature on deference uses the term ‘normative’ in ‘normative deference’, i.e. do they make the distinction between robust and formal normativity? This issue is rarely explicitly addressed, and so it is not always clear, but the tendency does seem to be to focus on robust normativity. For example, Fletcher (2016, endnote 25) expressly says that he uses “‘normative’ where others use ‘robustly normative’”, such that etiquette, club rules, and the like, are not normative.¹⁶²

There might be good reasons to focus on robust normativity. Perhaps one might think that formal normativity is too weak and not authoritative enough to be interesting. But whatever it is, there needs to be an explanation for why we should limit our investigation to one type of normativity and not the other. Moreover, a story is also needed for why we would say that deference about formally normative matters is not problematic (as is implied by exclusively targeting robust normativity). Anyone who wants to hold a domain-based asymmetry thesis, based on the (robust) normative versus the non-normative, owes us an explanation for all of this.

Nevertheless, however you conceive of the normative, the broad version of the asymmetry thesis fails too. Here is why. Take the view that there is a

¹⁶² Fletcher references Copp (2007, ch. 8), Enoch (2007), and McPherson (2011).

distinction between robust and formal normativity, such that deference about robustly normative matters is problematic and deference about formally normative matters and non-normative matters is not problematic. Take the example from above:

TRIP: Anna is a journalist who is preparing to go on a reporting trip to a dangerous and conflict-ridden area. She has to tell her family that she will be away but she really doesn't know what to tell them. If she tells them where exactly she's going and why, they will be extremely worried. On the other hand, she worries that by evading the questions she would be lying. She goes back and forth but cannot decide what the right thing to do is and eventually decides to ask a friend whose judgement she trusts. (Sliwa, 2012, p. 178)

This is an instance of moral deference and, the moral being part of the robustly normative, it is an instance of normative deference too. If we agree with Sliwa that this case is not intuitively problematic, and that there are others which are similar, then we have a type of counter-examples to the broad version of the asymmetry thesis. Not all cases of moral deference, though falling under the guise of (robust) normativity, will be problematic.

But here is an example of formally normative deference that does seem problematic: a case of etiquette deference.

POLITE. Jane keeps deferring about what is polite or rude (in her own society). She knows the concepts of politeness and rudeness, but has trouble figuring out what politeness asks of her in different situations: whether she ought to stay in a conversation that bores her to death; whether she should ask this acquaintance that she know is going through a rough time about their personal life; whether she should say the baby is cute when she thinks it is clearly not etc.

Plausibly, this example is about, or at least involves, norms of etiquette. My intuition is that such a case cannot be so easily dismissed as unproblematic — at least not if you compare it to cases of deference regarding what the club rules are or what cutlery to use at dinner. But, because it is a case of deference in the etiquette domain it would not count as the relevant kind of normative deference, i.e. not robustly normative. If the (robustly) normative and the non-normative

are to track the problematic-unproblematic deference distinction, then this does not align with our intuitions.

For those who do not subscribe to the formal-robust normativity distinction I have a question: why do certain cases of moral deference seem problematic and others do not? Why do certain cases of etiquette deference seem problematic and others do not? In other words, I am doubtful that a supporter of the broad version of the asymmetry thesis can explain why there is variation in the problematic character of different cases within a given domain. Since such an asymmetry is supposed to be domain-based, the intuitions should be that deference in some domains is problematic whereas deference in other domains is not problematic. But it turns out that it is not so clear-cut and that there is indeed variation within domains. My point is precisely that using domains as the basis of the asymmetry is too coarse-grained it cannot handle the fact that not all cases of deference in one domain will be the same.¹⁶³

What this discussion suggests is that it is fairly straightforward that we should not take the original Hopkins formulation of the asymmetry thesis literally. Discussing the asymmetry in terms of a difference in kind, interpreted here as a stark binary contrast between moral and non-moral deference, such that one is *always* problematic and the other one is *never* problematic, is implausible. Examples show that such an asymmetry thesis would be false. It would be false even in the broad version because that does not do justice to our intuitions either.

However, maybe the supporters of the domain-based asymmetry could say the following. Let us not say that the thesis is about a rigid contrast between moral or normative deference that is *always* problematic versus non-moral and non-normative deference that is *never* problematic; that might be too strong. We can say instead, like Howell (2014) does, that while there is something *generally* problematic with moral or normative testimony, there is nothing *generally* problematic with non-moral or non-normative deference. First of all, this would not be very helpful. As I have argued in chapter I, this is a very unclear formulation and as such it is not capable of sustaining a pessimist view — what

¹⁶³ It is difficult to find cases of non-normative deference that are problematic *per se*. Nevertheless, I do not think this shows that the broad version of the asymmetry thesis is right. For one, the Anna example shows that normative deference is not always problematic. And the discussion shows that there is variation within domains. I take it that together all these considerations give us reason to doubt even the broad version of the asymmetry thesis.

does ‘generally’ mean, after all, in terms of strength and scope? If we do not have answers to these questions, we will not be able to determine which cases will be problematic and which will not. Secondly, it is not sufficient to simply assert that normative deference¹⁶⁴ is problematic while its opposite is not. We still need an explanation for why that is the case. And the rival pessimist accounts do not satisfy this condition. What they do is give an argument for moral (or moral and aesthetic) deference pessimism and hope this extends to all normative domains. But that does not happen either.

Take the most developed pessimist position: the moral understanding argument. Can this argument apply to all normative deference? It does not seem so easy. Hills (2017b) attempts to show it extends to aesthetic deference, but the matter is not settled. Recall that her argument was centred on the notion of moral worth. Moral deference is bad because it precludes us from getting moral understanding, which is necessary for our actions to have (some higher degree of) moral worth (and virtue). This involved, among other things, doing the right thing for the right reasons. But is there such a thing as aesthetic worth, when it comes to actions that have to do with aesthetics? Moreover, are all our beliefs about aesthetics related to action in the same way that moral beliefs are (such that we have reason to worry about aesthetic deference precluding us from acting in an aesthetically worthy way)? Is there such a thing as aesthetic virtue?

Hills (2017b: 7-8) says that there is indeed aesthetic worth and we can first and foremost use it to evaluate actions that are part of creating a work of art. When an artist works to create some piece of art they can do so in response to aesthetic reasons or not. Only the former type of actions will have aesthetic worth. The actions of someone who produces art simply on the basis of deference will not have (high degrees of) aesthetic worth.¹⁶⁵ But, of course, one can defer about beliefs that will not be put into practice in the creation of a work of art. Since aesthetic *worth* applies to *actions*, what can be said about beliefs, to

¹⁶⁴ I target the broad version of the asymmetry thesis here simply because by accepting that aesthetic deference can be problematic too (which, I think, all pessimists would do), the narrow version immediately loses its plausibility.

¹⁶⁵ There is a further question whether artists make purely aesthetic judgements or artistic judgements. I interpret Hills as saying the former, because an artistic judgement is an impure aesthetic judgement, so to speak (i.e. it includes non-aesthetic elements related to, e.g., the history of art, the tradition one is working in, how the work will be received and recognised, the possible contribution made to the artworld). If the artist makes artistic rather than purely aesthetic decisions, that would pose further challenges for this view.

which this notion does not apply? If the challenge cannot be answered, then this account would leave out a significant number of cases of deference, and so the account would be underinclusive. Hills (2017b, p. 8, my italics) considers this issue:

So far I have discussed aesthetic judgements that influence the creation of art, but of course not all of them do. Some people produce or sustain the museums, concerts and so on that the rest of us enjoy. And the rest of us attend exhibitions, concerts, plays and cinemas, and recommend some of what we see to others. *This is not artistic creation but it is both active and responsive to aesthetic reasons. So it is (or rather, it can be) aesthetically worthy action.*

It is not clear why active participation and/or responsiveness to aesthetic reasons are/is sufficient for aesthetic worth. Perhaps one idea could be all aesthetic judgements are in some way practical and so are clear candidates for aesthetic worth. This is, however, a strong claim which needs support — this is a challenge that my own account will have to face and which I will expand on later.¹⁶⁶

Of course, this is early work from Hills and perhaps there could be a way to develop the view to overcome this challenge. So let us accept that this is possible, for the sake of argument. The danger is, then, that the notion of worth — and, indeed, that of virtue — will just apply to everything. There might be prudential worth and prudential virtue; etiquette worth and etiquette virtue; mathematical worth and mathematical virtue; perhaps even weather forecast worth and weather forecast virtue. However, for one, this seems implausible. Recall that (full) virtue here involves having the right beliefs, right non-cognitive attitudes, and doing the right actions, as well as the cognitive control abilities specific to Hills' view (i.e. the ability to follow give an explanation for a

¹⁶⁶ Hills (2017b, p. 9) discusses how cognitive attitudes can be responsive to aesthetic values. Aesthetic judgements formed in response to aesthetic reasons constitute appreciation (more precisely, the cognitive part of appreciation). Such an appreciative judgements mirror the content and the form of the value to which it is responsive. For example, saying that “*The Hunt in the Forest* is a good painting” in response to the relevant aesthetic reasons (rather than on deference) mirrors, first, the fact that the painting is actually good. Secondly, it also mirrors the fact that it is good in virtue of its colour and composition (which are exactly the reasons on which the belief is based). This kind of cognitive appreciative judgements paired with non-cognitive appreciative attitudes, and “together with appropriate action is what (full) aesthetic virtue requires” (Hills, 2017b, p. 10). This still does not settle the aesthetic worth challenge posed above and there are also questions about whether this claim is supposed to be about the necessary and/or sufficient conditions for virtue.

judgement, follow an explanation given by someone else etc.). It is not clear that this notion would apply well outside morality. To expand it all the way to, e.g., weather forecasting, is even less attractive. This would make the concepts of worth and virtue so diluted and different from their original usage that it would be difficult to actually treat them in the same way. Secondly, if we apply this to every single domain, then we will not be able to preserve the broad version of the asymmetry thesis. In fact, we might not be able to hold on to *any* kind of asymmetry: if all our beliefs and actions are evaluated by appeal to the ideal of worth and virtue, then any kind of deference will fall short of that and thus be problematic.¹⁶⁷

Things are not looking much better for the other pessimist views. The virtue account would face broadly the same problems. The authenticity account would likely not be general enough.¹⁶⁸ Take aesthetic authenticity. Perhaps this is most important for artists — such that their work is an expression of their true selves — but what about the rest of us? It is not entirely clear why the demands of authenticity would have a pull on our aesthetic lives, especially by contrast with our moral beliefs. Recall that the view uses the idea that moral beliefs are central to our identity. But are aesthetic beliefs also central to our identity? Mogensen (2015, p. 280) explains that in the 19th century, in the Western world, people started to put more emphasis on our subjective attitudes and sentiments in relation to aesthetic appreciation. That cultural turn has had an impact such that we now feel that our aesthetic judgements should be expressions of our sensibilities. However, this merely suggests that we came to think that it may be *important* to express our aesthetic sensibilities, but not that our aesthetic beliefs are central to our identity (at least not in the way moral beliefs are said to be on Mogensen’s view). And that is not sufficient because, as Mogensen worries himself, this might make the authenticity explanation vacuous:

We might worry that we’ve gained little explanatory insight unless we can also say why moral beliefs are

¹⁶⁷ Lord (2018, p. 87) doubts this on different grounds: “the aesthetic analogue of Hills’ view is far from obvious. It is not plausible that in order to be aesthetically virtuous one needs to understand why the aesthetic facts obtain. Such understanding is an intellectual good. However, when it comes to one’s aesthetic character, appreciation seems more central.”

¹⁶⁸ Skarsaune’s (2016) view is more difficult to expand beyond morality than Mogensen’s (2015). Recall that the former talks about the importance of authentic interactions, which depends on what people are involved and the gravity of the situation. So what I say here applies more to Mogensen’s account.

special in this regard. Why does authenticity make special demands of our moral beliefs? Why isn't our authenticity similarly at stake when it comes to beliefs about geography or zoology? (Mogensen, 2015, p. 277)

He responds to this question precisely by saying that when it comes to the moral case, one's moral beliefs are particularly central to their identity. This is a contentious claim in itself, but to hold it in relation to our aesthetic beliefs begs for even more support. If we do not have an argument for that, and the claim that a belief needs to be central to one's identity for the demands of authenticity to apply is not part of this position, then we have the vacuity worry indeed. The demands of authenticity would apply to all beliefs (where it can be applied), and so the account would overgeneralise. Why are there no demands of authenticity on, say, our beliefs about physics? For example, because of what I perceive, I may think that the Earth is flat —the authentic me has this belief. But, at least intuitively, I *should* defer to others when they tell me that the Earth is actually round. The demands of authenticity do not seem to apply here, even in a *pro tanto* way.

The acquaintance view shows some promise. Lord (2018) explicitly talks about acquaintance with *normative* properties. But the charge of underinclusivity transfers to deference beyond morality. Think about prudential deference, and the moral cases I discussed in chapter II as not being covered by Lord's account. That is, deference about how to weigh up different relevant reasons and deference about truths that are reached through theorising. In the first type of case, acquaintance is present, as the person involved is presently living a situation which requires making a decision, and so has direct access to the relevant reasons. As they reflect on the reasons that exist, they have trouble seeing which one prevails and so defer about that — Lord's account seems to predict that such deference should not be problematic, since the absence of acquaintance is the issue on this view. But, in this case, acquaintance is not absent — it is only that the person cannot weigh up the reasons. In the second type of cases, acquaintance is not even on the cards (given that theorising does not entail acquaintance anyway), and so such deference should also not be problematic. We can encounter the same kind of cases with respect to prudential matters. For example, I can be in the situation where I need to decide in what company to invest my money, to make sure that it benefits me the most. I can think of the

reasons why one company would be good, the reasons why another one also looks promising, and so on. I get acquainted with all the relevant normative features of the situation, and yet I end up uncertain on how to weigh them up in order to make a decision. If the lack of acquaintance is what would make prudential deference problematic, then if I were to defer in this case, it should be fine, according to this account, because I do have acquaintance. However, according to general intuitions, and to the supporters of the broad version of the asymmetry thesis, there is always something *pro tanto* bad about prudential (normative) deference. So the view does not test well here.

Moreover, many of our prudential beliefs also come from theorising about our well-being and interests. Think about general claims, such as “if A is better for you than B, *ceteris paribus*, you have more reason to do A” (Rodogno, 2015, p. 288). Presumably we come to such a conclusion by doing some armchair reflection, which does not include an element of acquaintance. Lord (2018, p. 93) admits some beliefs are gained like this, but he thinks it is rare that people get knowledge in this way, at least when it comes to moral knowledge.¹⁶⁹ That may well be, but the general intuitions point to the fact that deference about some such general or necessary normative truths is problematic. But, on the acquaintance view, it should not be problematic — there was no acquaintance in obtaining this belief by armchair reflection anyway, so nothing is lost by deferring instead.

Moreover, recall that this argument says acquaintance is important because the role it plays in us getting *appreciative knowledge*.¹⁷⁰ Moral and aesthetic appreciative knowledge is important (even central — Lord, 2018, p. 93) to our lives as moral and aesthetic agents. According to Lord (2018, p. 76), appreciative knowledge is knowledge that permits one to have all the appropriate conative and affective reactions to some properties. When it comes to morality, we can say this kind of appreciative knowledge would help us act in a morally worthy way and become more virtuous. In our aesthetic lives, we may think that

¹⁶⁹ My point stands regardless of whether or not Lord would want to say it is only rarely that we get new prudential knowledge through armchair reflection as well. Even in that situation (assuming it is true), there will still be a range of cases that his account will not cover, albeit a limited one.

¹⁷⁰ Also, perhaps Lord would want to say that acquaintance is *not* required for appreciative prudential knowledge, just like he says in the moral case. Then the issue with deference would be that it fails to give us appreciative knowledge.

appreciative knowledge is important because appreciation is essential when engaging with aesthetic objects. But does the claim that there is appreciative knowledge and it is central to our lives apply beyond moral and aesthetic knowledge, in accordance with the broad version of the asymmetry thesis? It seems to me the answer might lead to a dilemma.

Either this does not generalise, in which case the acquaintance view does not account for the broad asymmetry thesis after all, as it would only cover moral and aesthetic deference, and not all normative deference. For example, it is not at all clear that prudential appreciative knowledge is *central* for us as prudential agents; at least not in the same way (or perhaps not to the same degree). There does not seem to be a convincing analogue between moral and aesthetic appreciative knowledge, and prudential appreciative knowledge.

Or the claim does generalise, but then there seems to be no reason to not consider even more kinds of appreciative knowledge, thus going beyond the broad asymmetry thesis. For example, if there is such a thing as prudential appreciative knowledge, and it is central for us as prudential agents, then why would there not be other kinds of appreciative knowledge, that go beyond normative matters? Perhaps there is such a thing as mathematical appreciative knowledge or appreciative knowledge of chess. As I have said, appreciative knowledge here is knowledge that enables one to have the full range of appropriate conative and affective reactions (to some properties). So perhaps having mathematical or chess appreciative knowledge would mean using certain facts as reasons to be in awe with the genius of a mathematical proof or a chess strategy; to want to promote that proof or the strategy; to want to make sure that it is used by people, and so on. This could be important, or central, for us as mathematical agents or chess-playing agents (in the same way prudential knowledge would be central for us as prudential agents, whatever that would be). Of course, Lord might deny that such appreciative knowledge is central to our lives. However, then we would need an explanation for why moral and aesthetic knowledge are special in this way. Either way, the broad version of the asymmetry thesis would not be respected.

The lesson here is that neither the narrow nor the broad version of the asymmetry thesis is successful. The proposals that have been made, intentionally or not, are incomplete. They cannot deal with a variety of cases of deference. But

perhaps they could be extended in some way to avoid all these worries. However, I am doubtful that this is possible and, even if it is, talking about moral versus non-moral or normative versus non-normative deference does not appear fruitful to me. There does not seem to be a principled way to draw the asymmetry between problematic and unproblematic deference based on domains of discourse, given the variation that can exist within a given domain. Moreover, it poses other issues too. One is that the concept of a ‘domain’ is murky in itself and there are no strict delineations between domains. There are cases where it is unclear whether the testimony belongs to one domain or another. For example, the moral and the prudential can be intertwined. Or many instances of political or legal deference may exist at the border between the political, respectively the legal, and the moral or even the prudential area. The moral and the non-moral can be entangled too.

Thus I think we should abandon the quest of trying to draw the asymmetry of deference according to domains. However, this does not mean that there is no asymmetry to pursue. As I have shown in the beginning of the chapter, our practices and intuitions strongly suggest that there is one. So we need to find out whether we can make sense of it in a systematic way. That is what I set out to do in the next section.

3. A New Basis for the Asymmetry Thesis¹⁷¹

Let me take stock: widespread intuitions and practices suggest that not all deference is the same. We seem to find some cases more acceptable than others. Looking at what domains those matters belong to did not help. It gave wrong predictions and it did not account for our intuitions. What we need is a criterion that can help us distinguish between problematic and unproblematic deference in a way that is not restricted to domains. So now I propose a new explanation: I expand the practical deliberation argument I have put forward in the previous chapter beyond morality. That is, I will use interference with practical deliberation as a criterion to draw a principled contrast between

¹⁷¹ I keep referring to the *asymmetry* thesis, but my account will reveal more of a contrast rather than an asymmetry. However, for the sake of continuity and to avoid introducing unnecessary new terminology, I will keep using the term ‘asymmetry’, along with ‘contrast’.

problematic and unproblematic instances of deference. The main claim will be that some pattern of deference is *pro tanto* bad because, and insofar as, it interferes with the exercise and the development of our valuable capacity for practical deliberation. Unlike the rival pessimist views, since the jurisdiction of practical deliberation extends beyond morality, my practical deliberation account has quite an easier time applying to other kinds of deference. This is, I believe, one of the biggest virtues of my position. It is the most extensionally adequate account available, being best equipped to cover the widest range cases of deference and to systematically distinguish between the problematic and unproblematic ones in a way that aligns with our intuitions.

To briefly remind my reader, I take the capacity for practical deliberation to be what enables us to engage in processes of practical deliberation in order to work out what to do, in response to reasons. It typically includes reflection on, or imagination of, possible lines of action, taking into account different considerations and checking their relevance, weighing them up. This kind of process is carried out with the aim of arriving at a decision, intention, or action. I claim that the exercise and the development of the capacity for practical deliberation are both instrumentally and finally valuable.¹⁷² They are instrumentally valuable because practical deliberation can help us solve our practical questions and decide what to do, guided by reasons. In this way, we (and our future selves) can deal with complex situations on our own, guided by reasons-responsiveness, without needing to worry about having an expert next to us at all times from whom to take testimony. The exercise and the development of the capacity are also (extrinsically) finally valuable because in deliberating about our decisions and questions we shape and specify our values and ends, and as such we form our normative view of the world. Because we do not have ready-made answers for all our practical questions and we do not have the exact specifications of all our normative commitments, we need a way to do that. We need to shape them such that we shape our lives. This is precisely what practical deliberation does. It is precisely the procedural skeleton on which this kind of life is built. This kind of life could not be constructed without practical

¹⁷² As I have said in the previous chapter, I take this final value to be conditional on the kinds of creatures that we are. Had we been different, the way we lived our lives and what we valued might have been different.

deliberation. In this sense, practical deliberation and this kind of life are valuable as ends.

The problem is, I hold, that, in certain cases, recurrent deference interferes with this valuable capacity. More precisely, it precludes the agent from practising their practical deliberation and so it interferes with its exercise and its development. The interference happens in the following way: instead of deliberating, the agent simply defers to someone else about what the answer to the question they are struggling with is. Deliberation is replaced by deference (or some of the steps it involves are replaced). In this sense, the exercise of the capacity for practical deliberation is interfered with — instead of exercising our practical deliberation, we defer. Moreover, I hold that its development is circumvented too. Like most other capacities, the capacity for practical deliberation needs to be practised to improve. But certain cases of deference eliminate the possibility of practice precisely by providing a shortcut to the answer. In this sense, the development of the capacity for practical deliberation is interfered with — we do not get the chance to exercise it and without exercise we cannot improve it. So, since it is *pro tanto* bad to interfere with something of value, as long as some pattern of deference interferes with the valuable capacity for practical deliberation, such pattern of deference is *pro tanto* bad.

The claim I want to make here is that a view along these lines can provide a principled and extensionally adequate distinction between problematic and non-problematic cases of deference. This is because the jurisdiction of practical deliberation is broad: we use this capacity in relation to a variety of non-moral, normative, and non-normative issues. The capacity for practical deliberation is not domain-specific. For example, moral deliberation is not different in kind from aesthetic deliberation — they have the same procedural skeleton. My account then predicts that any kind of deference that interferes with the exercise and the development of our capacity for practical deliberation will be problematic, i.e. *pro tanto* bad, to that extent. In the next section I will put my account to the test, by considering various cases of seemingly problematic and non-problematic deference and show that the predictions it gives are in accordance with our intuitions.

4. The Asymmetry Thesis Revisited

4.1. Theoretical Deference

Here is my position: there is indeed an asymmetry (or better: a contrast) between different cases of deference such that some are problematic and some are not. The contrast is between deference which interferes with our capacity for practical deliberation and deference which does not. The former is problematic, the latter is not. So when confronted with some pattern of deference the question to ask is: does it interfere with our capacity for practical deliberation?

Start with the most widespread cases, namely the ubiquitous reliance on testimony about everyday (and not so everyday) theoretical matters. These are the kind of cases which we generally deem unproblematic: deference to news programmes and newspapers about current events; to professors about physics, chemistry, or history facts; to weather forecast people about the weather; to our accountants about how much we can spend and how much to save up etc. We recurrently defer about such matters and it does not seem suspicious. The defenders of the old asymmetry thesis would say this is so because it is non-moral or non-normative deference. Of course, we have seen that this explanation is too coarse-grained and the following discussion will reinforce that.

My account can neatly explain why this kind of deference is generally unproblematic. The matters under consideration do not typically fall under the jurisdiction of practical deliberation and thus the capacity for practical deliberation is not interfered with when we defer about them. Rather, they are theoretical, and we would not even use practical deliberation to think about them. Moreover, we often could not deliberate our way to the answers on our own anyway. For example, no matter how hard I try on my own I could not deliberate my way to knowing how certain molecules interact. I need to learn about it by deferring to the experts. I also cannot deliberate my way to tomorrow's weather forecast. So, of course it is not problematic to defer about these theoretical issues.

However, this may not hold for all cases of deference about theoretical matters. There are some special instances where theoretical deference which

seem problematic. First, at least some cases where we defer on yes or no questions, with no additional explanation, may be suspicious. For example, you ask an expert ‘Is string theory true and does it describe reality well?’ and they just say yes. You adopting the belief just on this unexplained answer might well trigger negative intuitions. But my view could not explain why such a case might be problematic because no practical deliberation was involved. If deference replaced anything it would be theoretical deliberation. Theoretical deliberation, or theoretical reasoning, is taken here to be another species of deliberation, just like practical deliberation is. Theoretical deliberation is about what to believe, in the sense that it is a process of weighing up options and evidence, that results in a belief, “or a new credence in a possible state of affairs, or a sincere thought about the truth of a proposition, or something similar” (Arpaly and Schroeder, 2012, p. 210). Theoretical deliberation is typically not concerned with what to do, either in issue or in subject matter. In this sense, no interference with *practical* deliberation occurred in the case above, so such deference should not be problematic, under my account.

At the same time, other plausible explanations present themselves. In deference to yes or no questions type of situations, I suspect the issue has to do with the lack of any additional explanation. As such, the intuitively problematic character of this kind of deference might be explained not as *pro tanto* non-epistemic, evaluative, badness of the sort discussed above; rather, it would be in terms of a violation of a different kind of norm, such as an epistemic norm, a norm of assertion, a conversational norm, or some other type of social norm. For example, one might say that deference here is epistemically risky: the issue is controversial and there is a lot of reasonable disagreement about it in the scientific world. Or, alternatively, one might say that, conversationally, it is strange that someone just answers ‘yes’ or ‘no’ when they are asked about a complex matter (they might be breaking the Gricean maxim of quantity, as they have not offered enough information, or the maxim of manner, as they have been a little too obscure). Equally, it might also seem strange that someone does not ask for details when they want to know the answer to a complex question.

Nevertheless, I do not think such cases represent a counterexample to my view. My account does not present itself as being the sole explanation for every case of problematic deference that exists. On the contrary, it is compatible

with there being other explanations for various examples of problematic deference, maybe even some that are specific to theoretical deference. However, I will not attempt to find out which such norm is broken here. It suffices to say that if it is plausible that these cases of deference are problematic in virtue of violating some such other norms, then they are not counterexamples to my account, which ranges over cases where deference is non-epistemically, evaluatively, *pro tanto* problematic in itself. Nor will such cases make my account uninteresting, since they hardly will exhaust the problematic cases of deference.

Second, there are further particular cases of theoretical deference that seem problematic, namely those involving specialists deferring in their areas of expertise. Again, my account might not be able to explain them because there is no interference with *practical* deliberation going on. Think about how strange it would be if a physicist, who is not in training, recurrently defers about matters in physics, in particular from the area in which she works. Or imagine a philosopher who too often defers to fellow philosophers.

But the negative intuitions would actually hold for cases of experts who defer on matters covered by the jurisdiction of practical deliberation too. Take the following situation:

Experts: Consider a brain surgeon whose knowledge of the whereabouts of various parts of the brain and how to manipulate them with disturbingly small, sharp instruments comes from testimony (what we might call brain-surgery testimony). Let's imagine for the time being that the surgeon is really very good with the scalpel. He can do everything that needs to be done so long as he is told *exactly* what to do by Cyrano de Braniac, who whispers to the surgeon each move he ought to make. What would we say about this surgeon? (Groll and Decker, 2014, p. 61)

It also seems bad that the brain surgeon can perform very good surgeries only if he is told exactly what to do and is constantly deferring to another actually brilliant surgeon. Assuming he is not in training, this case seems problematic. It would seem equally strange if an applied physicist would defer about how to set

up experiments or if a judge would defer about what decisions to make in court¹⁷³ (again, all of this holds if these people are not in training and they cannot be expected to know these things).

Luckily, I think these cases also fall out of the category I investigate: they are not problematic in themselves. Generally, we expect that people who undertake a certain specialised role do what the job description requires of them. If one is to fulfil a certain role, it seems they ought to do what their role entails on their own rather than to constantly defer. Deference seems problematic in such cases because it reveals that they are not capable of doing our job and, perhaps, that they should step out. This explanation holds independently of our views on deference as it appears to refer to some professional norms, the violation of which triggers our negative intuitions.¹⁷⁴ So, again, these cases do not represent counterexamples to my view.

4.2. Prudential Deference

Now to return to the jurisdiction of practical deliberation, consider some cases of deference that are covered by it: cases of prudential deference. My account predicts, mapping onto our intuitions, that recurrent prudential deference is indeed problematic because it interferes with the capacity for practical deliberation.

The prudential is concerned with what is good for us, what is in our interest, often closely related to the notion of well-being (Taylor, 2013; Tiberius, 2015). For example, prudential value has been defined as:

One of the things typically considered when figuring out what options, experiences, pursuits, or kinds of lives to pursue or choose. What's best for me or in my best interest, the life of pleasure or the pursuit of knowledge?" (Rodgno, 2015, p. 288)

¹⁷³ Hills (2009, p. 110) uses the example of the incompetent judge, but her explanation for why their deference is problematic is the lack of moral understanding and reasons-responsiveness. She takes this case to be problematic *per se*. I am not sure I agree — our intuitions seem overdetermined here so it is at least not fruitful to use it.

¹⁷⁴ This 'doing one's job' explanation has been discussed in the literature (Groll and Decker, 2014; Howell, 2014), but the focus was on why it cannot generalise to morality.

So here is a case of prudential deference: I defer to my friend about whether I ought to make a risky financial move, like buying some expensive but potentially rewarding stocks that have been a little unpredictable lately. We can also imagine cases of recurrent prudential deference, where I defer about things like what career path to choose, with whom to have relationships, or to what I should dedicate my spare time.

Prudential deference has received much less discussion than moral deference.¹⁷⁵ But it seems to trigger negative intuitions in the same kind of way as moral deference. One *prima facie* explanation for this is simply that no one else knows what is better or best for a person, but that person; or, at least, no one should get to decide that for them. But this is not always true. Sometimes those outside the situations we are in are able to better comprehend and evaluate what we should do. Our judgements might be clouded, we might be over-invested, have false beliefs etc. Other people may know what is better or best for us. Of course, this does not mean that they *should* get to decide for us; maybe they just should not. But even if this restriction holds, it would not explain why prudential deference seems problematic. Especially because we could voluntarily give up opportunities to decide for ourselves and let others do it for us, and prudential deference would still seem problematic. Additionally, if we think that letting others decide for us is the issue, then prudential deference would not be problematic in itself, but only because it would violate an independent norm and thus it would have nothing to do with deference. This would then not be a restriction on deference, but rather on how to live with each other.

According to my account, there is a better explanation for our negative intuitions. For one, the jurisdiction of practical deliberation naturally extends to the prudential. The prudential can get pretty close to the moral (or even overlap with it), which is paradigmatically practical, so it should not be controversial that we will need to do quite a bit of practical deliberation to be able to figure out how to go about our prudential problems and decide what is best for our well-being. We will need to do some work to identify the prudential reasons that apply to us, to define and specify what well-being is for us and what it includes, what has prudential value, how to specify our prudential values, how they relate to our other values and how they weigh up against them, how to obtain what is

¹⁷⁵ There are some exceptions, e.g. Driver (2006); Fletcher (2016).

prudentially best for us. Just like in the moral case, our prudential commitments will need qualifications and sometimes we will need practical deliberation for that. But this is precisely what deference would prevent us from doing by repetitively replacing the exercise of practical deliberation with testimony. In this way, prudential deference interferes with our capacity for practical deliberation. That is why my view deems prudential deference *pro tanto* bad, confirming our negative intuitions about it. Of course, this is what the broad version of the old asymmetry thesis would say as well (though not the narrow one) — that prudential deference is problematic. Here, our accounts converge because prudential deference simply happens to be normative deference. As shown above, however, it is not clear whether any of the other existing pessimist accounts would be sufficiently well-equipped to actually cover prudential deference. Nevertheless, the fact that my argument gets it right in more cases than the old asymmetry thesis is evidence enough that it does better overall.

4.3. Etiquette Deference

Now I want to tackle a topic that has not been discussed in the deference literature and make the case that it is worth considering: etiquette deference. Etiquette has not even been mentioned in the debate. It is automatically presumed to be of no interest because of cases like the following:

TRAVEL. Steve goes travelling for a year in Asia, where he has never been before. He wants to be polite and considerate to his hosts, but he has not had time to do his research on the local etiquette of dining. So when he gets there, he asks people what the customs are: how he should serve himself, what kind of instruments he should use and how, whether he should finish the food on his plate, whether he should make sounds when he is eating, and so on. Based on the locals' testimony he forms many beliefs on the etiquette of dining and he acts accordingly, deferring in each region and country, for the whole duration of his trip.

We seem to think there is nothing bad about Steve's deference about the different Asian dining customs. We may even generally have this idea that etiquette is only, or mainly, about what cutlery to use first at an elegant dinner or what to wear to the country club — it concerns the sort of formal and non-

authoritative normativity that is not so important. I want to try now and push back on this view about etiquette: there are reasons to believe that not all etiquette deference is so easily dismissed. That is, that there are divisions within etiquette taken as a domain, such that some patterns of recurrent etiquette deference are problematic and some are not. This further gives support to the idea that a domain-based asymmetry cannot adequately map onto the variety of some given problematic and unproblematic cases.

Etiquette has been defined as “a body of rules concerning proper conduct in everyday encounters” (Yeung cited in Burns Coleman, 2013, p. 69). It is usually discussed as a code of customs that governs the behaviour of people in general or particular groups of people. Roughly, it is a set of implicit social rules that apply in different situations, to which someone can voluntarily subscribe. There are rules of etiquette regarding many things, from how to eat and how to dress to how to respond to correspondence and how to greet people; there are etiquette rules in many domains, from philosophy to football and from art to medicine etc. For example, it is considered part of the etiquette of different sports that the opponents shake hands at the end of a game or competition. Or, in the UK, it is customary that, at a wedding, the father of the bride gives the first speech. *Prima facie*, it does not seem problematic to defer about these types of things, just like in TRAVEL.

I hold that the explanation for this is that we cannot deliberate our way to this information. Although these are matters which fall under the jurisdiction of practical deliberation, they concern certain rules and norms that are conventionally fixed and which we have little reason to know if we do not play that game, so to speak. So, it is okay to defer about them; we could not get the answers otherwise. That is why our capacity for practical deliberation is not interfered with in such cases: we could not deliberate our way to the answers anyway.

But, not all instances of etiquette deference are like this. Whether some pattern of recurrent deference is problematic varies across different cases. This is because, as Buss (1999, pp. 795-796) puts it: “the most important lessons in manners are the lessons in how to avoid being discourteous, impolite, rude, inconsiderate, offensive, insulting”. Not all etiquette norms are about what to wear to the country club. Some etiquette matters will engage our practical

deliberation and will have an influence on the way we shape our lives and how we want to be as people. Take the POLITE scenario from above, where Jane keeps deferring about what is polite or rude. She knows what politeness and rudeness mean, but she does not always know what the polite thing to do is. Should she stay in a boring conversation? Can she ask this acquaintance about her personal life, even though she knows something rough is going on with him at the moment? Should she say the baby is cute even though she really does not think so? This example is about social norms. That is, about norms of politeness — etiquette norms. But it does seem to be just like Steve recurrently deferring about the different Asian dining etiquette norms. I would say that it is different and, while Steve's deference is not problematic, Jane's deference is.

One reason for why this is the case could be that, as some suggested, there might be a close connection between etiquette and morality.¹⁷⁶ The views vary regarding how deep this relation is, but I find it plausible that some etiquette matters (can) make reference to some moral or even prudential values. If we think that moral and prudential deference are problematic, then that could explain why some etiquette deference is problematic, namely that which engages moral and prudential values. But I do not want to restrict my claim to this. Whether or not etiquette indeed makes reference to these other areas, some issues nevertheless engage our practical deliberation and have an impact on how we develop it and then further on how we live with other people and what kind of people we will be. Think about how we sanction people who recurrently violate etiquette norms; we dislike rude and offensive people, consider them disrespectful, and think their behaviour shows something about their worldview and who they are, even though not necessarily their moral views. Some etiquette matters will require of us to deliberate: to try and find the relevant reasons (is the fact that not shaking hands with women makes them uncomfortable a reason to start doing so?), to see how considerations weigh up (what is more important: my desire to stop looking at baby photos or my friend's happiness that he is able to share them with me?), to take perspectives and imagine different scenarios (how does my colleague feel if I do not even ask her how she is? Will she be colder to me if I do not?) etc. Plausibly, some etiquette deliberation could develop our capacity for practical deliberation and develop important

¹⁷⁶ See Martin (1993), and Olberding's discussions (2015) that make reference to Confucianism.

sensibilities. But, again, that is precisely what deference stops us from doing, namely we do not get to exercise doing all of this. Insofar as such deference compromises our practical deliberation, it is *pro tanto* bad.

As mentioned, the defenders of the old asymmetry thesis do not discuss etiquette deference. They might not be interested in it because not everyone believes that it is the relevant kind of normative deference. It is seen to involve merely formal or non-robust normativity. As said above, this sort of view leaves much to be desired. More importantly, it begs for an explanation of the distinction and of why only one type of normativity is problematic. Regardless of that, I hope to have shown that not all etiquette deference is the same and that it does not matter whether we think it is formally or robustly normative. Some patterns of etiquette deference will trigger negative intuitions and will interfere with our practical deliberation. My account can deal with this type of deference. The old asymmetry thesis could not, even if it accepted that etiquette is the relevant kind of deference. This is because it is domain-based and cannot account for divisions within one domain.¹⁷⁷

5. Objections: Aesthetic Deference

To take stock of what I have done so far: by looking at different cases of deference, I have reinforced the idea that a domain-based asymmetry thesis is wrong, as the contrast between problematic and unproblematic deference cuts across domains. Then I have presented my practical deliberation account and argued that it gives the right results in accordance with our intuitions when tested against various cases of deference. However, I have neglected one of the most discussed kinds of deference: aesthetic deference.

¹⁷⁷ A similar argument can be made for legal deference. The legal is also considered formally rather than robustly normative so, again, it is ignored by the debate. But I think some cases of legal deference can trigger negative intuitions. For example, deference about things such as how we should specify different legal values like justice and freedom, what we should include as human rights, what the best laws on punishment should look like, how to vote in a referendum that is about eliminating or introducing a certain law (e.g. marriage equality laws), what laws are good etc. My view has a way of explaining why such a pattern of deference would be problematic: it would interfere with practical deliberation (even though deferring about, e.g., what a certain law says, would not). The old asymmetry thesis, especially if it rejects the legal as merely formally normative, would not be able to account for these examples.

Consider the following example: my friend and I are looking through a book and we see Van Gogh's 'Starry Night'. My friend says: 'wow, this is a beautiful painting'. I look, but I do not really see what she means, I do not see the beauty. Actually, I do not quite know what to think of it. But I know my friend has a good eye for art. So I just adopt her belief and go on thinking that 'Starry Night' is a beautiful painting. Now imagine that I do this kind of thing multiple times. I defer to my friend about the beauty of other paintings; I also defer about the aesthetic virtues of different sculptures, of certain pieces of music, of the ballet shows we see together etc. In other words, I recurrently outsource my aesthetic beliefs from her. Many people find this kind of aesthetic deference problematic. (Hopkins, 2011; Lord, 2018; Nguyen, 2017; Whiting, 2015).¹⁷⁸

Aesthetic deference poses an objection to my account of when and why deference is bad. My view is that interference with practical deliberation makes deference bad. But we do not usually think that the aesthetic — or at least a large part of it — falls under the jurisdiction of practical deliberation. In an intuitive sense, it is very different from the moral or the prudential, and it is not generally seen as answering practical questions. As aesthetic agents, we mainly form aesthetic *judgements* and we *appreciate* aesthetic objects. So if the aesthetic is not even practical how can it interfere with practical deliberation? How can my account explain the badness of aesthetic deference?

In what follows I argue that my account actually can explain the badness of *some* cases of aesthetic deference, namely those which involve interference with our practical deliberation. It is true that the view is limited in this sense, but then no position can do better, as I will show towards the end. I start by challenging the claim that aesthetic matters are purely theoretical and argue that at least some of them are actually practical, either in their subject matter or in their upshot. That will then show that (some) aesthetic deliberation can be seen as practical deliberation. This in turn allows me to argue that aesthetic deference can interfere with practical deliberation in the same way that moral or prudential deference does.

¹⁷⁸ Among the philosophers who do not think aesthetic deference is problematic are Laetz (2008); Meskin (2004); Robson (2018).

To begin, note that there are clear cases where aesthetic judgements can be seen as practical, namely when they are part of a creative process of an artist. Someone might want to create a new painting and they deliberate about how to do that. They will undergo a process of practical deliberation. They would reflect on what colours to choose, how to combine them, what kind of shapes would work better, what sort of structure would transmit the message in the way they want. The judgements they will make will give them reasons for actions: to draw a certain line, to use red, and so on. This seems, at first sight, a minor point. But if we think about the fact that many of us use our aesthetic judgements to make our own creations maybe every day it does not such a limited range of cases (although perhaps what regular people do create might not be considered art as such): when we write poems or paint to relax, when we decorate our houses, when we look after our fashion style. The aesthetic is no stranger to the practical.¹⁷⁹

Secondly, take a recent argument according to which art criticism can be seen as practical rather than theoretical reasoning because it provides reasons for action rather than belief, i.e. reasons to engage with artworks in certain ways (A. Cross, 2017). Of course, the artistic and the aesthetic are (at least in principle) different and separable (Best, 1982; Kulka, 1981), but we can use this idea about art criticism to make an analogous point about the aesthetic. Cross's argument is the following: art critics who characterize and explain works of art for us do so by drawing our attention to certain aspects which they want to count in favour of not merely having certain beliefs about the artworks but also of certain ways of looking, listening, reflecting, and generally engaging with the aesthetic objects. In short, they give us practical reasons for acting and reacting in certain ways:

For example, telling me that the figures in the Botticelli are overly delicate might count in favour of paying close attention to those aspects of the painting, or comparing Botticelli's depictions to other painters in the Florentine school. (A. Cross, 2017, p. 307)

¹⁷⁹ Hills (2017b) makes the point about the practical aesthetic judgements of artists. Ridley (2015) argues that moral and aesthetic competences *qua* practical competences are structurally isomorphic.

I think it is plausible that this can extend beyond art criticism, to the aesthetic. Aesthetic utterances in general, such as ‘The Dance of the Cygnets is graceful’ or ‘Sunsets at sea are beautiful’ could be interpreted as carrying with them invitations to engage with these aesthetic objects. They could suggest that we should see this particular dance and pay attention to see its elements of gracefulness or that we should make the time to admire a sunset when we are near a sea; that we should appreciate these aesthetic objects.

There is a question about whether or not this is a conceptual claim, i.e. whether aesthetic concepts have it in their content or meaning that they are evaluative and thus that they recommend certain ways of acting and reacting. This is a complex topic that I cannot explore and do justice to here.¹⁸⁰ All I can say is that *prima facie* our aesthetic utterances do seem to carry certain invitations with them. As aesthetic agents, we share aesthetic information in our aesthetic community, and one way to do that is via our aesthetic statements. It is very plausible that aesthetic utterances (and the judgements they express) are evaluative, and that is all I need — what is going on at a more fundamental level, whether such judgements are evaluative in virtue of their meaning (content) or their use, is too difficult to settle here, and I do not take it as necessary to do so. If this claim has any *prima facie* plausibility, then this is another pillar of support for the idea that aesthetic judgements can be practical.

Finally, there are reasons to believe that we actually engage our practical deliberation when dealing with some aesthetic matters. This draws on the long-standing debate on the relation between the aesthetic and the moral (e.g. Carroll, 2000; Gaut, 2001). The question that is at the centre of the debate is whether the moral aspects of a work of art matter aesthetically. But everyone pretty much agrees that a lot of aesthetic objects and works of art do make reference to, or even revolve around, various moral elements. Kieran (2006, pp. 129-130) gives some examples:

Consider the differences in moral outlook between Edith Wharton, Hardy, D. H. Lawrence, Francis Bacon, Hemingway, Brett Easton Ellis, Martin Amis, Philip Roth, Hubert Selby Jr. or Carol Shields. What’s represented as worthwhile or being appropriate objects

¹⁸⁰ See Väyrynen (2016) for an overview of the related literature on thick concepts.

of pride, indignation, sympathy, shame or guilt vary immensely in ways that are at odds with each other.

This suggests that insofar as different aesthetic matters intertwine with different moral matters (and probably prudential matters), some aesthetic judgements will be practical in subject. So, engaging with such art would also engage our practical deliberation. Figuring out what messages different works of art are supposed to transmit, trying to interpret the moral lesson they convey, engaging with and evaluating the moral elements they present, reflecting on the moral reasons they discuss, imagining how they relate to our lives, taking the perspective of one of the aesthetic actors (be it author, subject, a character) etc. are just ways of doing practical deliberation. They might not always result in intentions or actions, but they do seem to engage us in the same way that practical deliberation in real life does. By doing this kind of aesthetic deliberation we actually develop the sensitivity and reasons-responsiveness that will be valuable across different areas, as we will employ them regarding different matters. And not only this: our aesthetic views contribute to how we see the world, how we lead our lives, and how we construct our normative outlook in general. Whether we find aesthetically pleasing things important, what kind of aesthetic objects we like, what place they have in our lives, how much time we devote to aesthetically charged activities, what we learn from aesthetic objects — all of these influence who we choose to be and how we choose to live our lives.

Perhaps each of the points presented above does not make a strong case on its own. But put together they represent at least some evidence that it is not wildly implausible that some aesthetic deliberation and judgements are practical. And if some aesthetic matters are practical and engage our practical deliberation, then some cases of aesthetic deference have the same features as the other cases of deference discussed above. So, for example, deferring about the interpretations or the value of some works of art rather than reflecting on them on our own would be similar to some cases of, say, moral deference, in that we would be giving up opportunities for practical deliberation. Thus, I think my practical deliberation view can apply to such cases. That is, in some situations, recurrent aesthetic deference can interfere with our practical deliberation in the same way, and insofar as it does, it is *pro tanto* bad because of it.

However, there might be a further problem here: the view does not seem to generalise to the full range of intuitively problematic cases of aesthetic deference. Some aesthetic judgements will not engage practical deliberation but deferring about them will still be fishy. Recall NEFERTITI, where Hannah formed the belief that Nefertiti's bust is beautiful just because Clara said so. Some rival accounts might be better equipped to deal with such cases. For example, remember Lord's (2018) account, discussed in chapter II. One might say that many aesthetic judgements are subject to an acquaintance requirement. On one such view it is argued that to have *appreciative* aesthetic knowledge one needs to be acquainted with the aesthetically relevant properties of an aesthetic object.¹⁸¹ Acquaintance is rationally important because it puts us in a position to “manifest knowledge about how to use those facts [the relevant normative properties] as reasons for appreciation.” Acquaintance is *required* for appreciative aesthetic knowledge, which is “the sort of knowledge that allows one to fittingly have the full range of affective and conative reactions” (Lord, 2018, p. 77). Deference does not enable us to obtain appreciative knowledge “because it does not put us in a position to be acquainted with the specific ways in which the aesthetic features are realized”. But appreciative knowledge is very important to our aesthetic lives and this is what makes aesthetic deference problematic.

I note three main problems with this view. One is that it does not generalise to the full range of problematic cases of deference either. Think about the ‘Starry Night’ case from above, where I defer to my friend about the beauty of the painting because, even though I am looking at it, I just do not see what that is beautiful. Here, although I have direct epistemic access to the aesthetic object, I cannot grasp its aesthetic properties (presumably because of some fault of my own). Acquaintance is not available to me here. But why would it matter whether I defer or not then? Deference is said to be generally amiss because it does not give us acquaintance and appreciative knowledge, but here both are out of reach anyway. I am not losing anything by deferring (at least not anything that has to do with the rational advantages of acquaintance). Deference does not put me in a worse position with respect to acquaintance or appreciative knowledge,

¹⁸¹ There are variations of this view. Some people think that *to make an aesthetic judgement* one needs to be acquainted with the aesthetic *object* (see Wollheim's original Acquaintance Principle, in Robson, 2012; Gorodeisky, 2010).

so it should not matter — from this perspective— whether or not I defer. But the intuition seems to be that me deferring to my friend about the beauty of ‘Starry Night’ is indeed problematic. This account cannot explain why.¹⁸²

A related problem is that it seems that can one gain appreciative knowledge if one receives aesthetic testimony that enables them to imagine the aesthetic object in enough detail. This is because Lord (2018, p. 91) explicitly says that we can gain appreciative knowledge through imagination. So, if in the NEFERTITI example above Clara describes Nefertiti’s bust to Hanna in a very detailed way, then Hanna might be able to imagine Nefertiti’s bust in a detailed way. And if imagination is enough for acquaintance, then it seems she becomes acquainted with the particular features of Nefertiti’s bust.¹⁸³ This would make acquaintance compatible with at least some cases of deference. But this would not seem to make a difference to how problematic the case is.¹⁸⁴ Now, to be fair, Lord focuses on aesthetic deference about pure assertions, as he calls it, where no explanation is given in addition to the aesthetic judgement. But then this means that his view might not cover the case above, when additional explanation is given and yet it still seems like a problematic case of deference. Thus his account seems to not generalise properly either.

Finally, to reiterate, another worry about this view is that it does not generate a plausible asymmetry thesis. Lord explicitly supports the broad version of the thesis, namely that normative deference is problematic and non-normative deference is not. However, I have shown in previous sections why we have reasons to reject this view. But even if we were to accept it, we may worry that

¹⁸² In a response to Hopkins (2011), Lord mentions that in cases like ‘Starry Night’, where one is uncertain about what to believe, there might be *epistemic* reasons not to defer. If it rational for me to be uncertain after seeing the photo of the painting, then I have epistemic reasons not to form the belief that it is a beautiful painting. The evidence I got from looking at the painting “will generally be of higher quality” (Lord, 2018, p. 80) than the evidence that my friend gives me through her testimony. The former will outweigh the latter and so it is epistemically impermissible for me to defer. However, I am not sure that it is true that my perceptual evidence is necessarily of a higher quality than my friend’s testimony. Perhaps (or we can stipulate that) I am just bad at judging art, and I do not have the eye for it, while my friend does. And I may want to acquire some aesthetic knowledge. It does not seem to be clearly epistemically impermissible to defer here and so I do not think this adequately explains the negative intuitions about this case.

¹⁸³ There is a separate worry regarding whether imagination can indeed make us acquainted with the properties of an object and whether it is sufficient to ground the full range of affective and conative attitudes that Lord discusses.

¹⁸⁴ If asked how she knows that Nefertiti’s bust is beautiful, Hanna would have to say: ‘I imagined it, based on my friend’s testimony’. But that seems strange and I am not sure we would take this to be on a par with actually seeing the bust. This is part of what needs to be worked out in full detail in a theory that takes acquaintance to be possible via imagination.

his account cannot even support that version of the asymmetry thesis. Besides the challenges I put forward in section before, we can also argue that the assumptions he relies on in order to extend his view beyond aesthetic deference are too controversial. For example, when discussing moral deference, Lord argues that we can get acquainted with moral facts through our conative and affective states. He claims that such states represent perceptions of the normative. He also thinks that we become acquainted with contingent moral truths through imagination. It is clear that these are not uncontroversial matters and it is not easy to see whether this view could expand to any other kinds of deference (without ending up over-expanding too, see section 2 above, pp. 159-160). So, I do not think the acquaintance view can sustain any kind of plausible asymmetry thesis, even if it were to do better than my position regarding aesthetic deference.

This means that we are left in a theoretically challenging position: there is no account that can give a unified explanation for all the problematic cases of aesthetic deference. As such, we will have to decide which view does a better job at explaining what is non-epistemically wrong with deference overall. I believe this is the major advantage of my view: my practical deliberation account is preferable to its rivals because it is the most comprehensive one on the market. Firstly, it is the only pessimist account developed to cover more than one or two types of deference — as the existing literature covers, at best, moral, sometimes aesthetic, and very rarely prudential deference. In doing so it also does not rely on very controversial assumptions, as practical deliberation seems to naturally cover a diverse range of cases of deference while also leaving room for variation. Secondly, it is the only view that systematically analyses the asymmetry thesis and offers a principled way of drawing it that is in broad accordance with our intuitions and practices. Thirdly, it is as close as we can get to a unified account as it explores the way different cases of deference pattern with one another, and it gets most cases right than any other view around. This is at least *pro tanto* better than just having dissipated explanations in different cases for what seems like a generic phenomenon.

6. Concluding remarks

In this chapter I have argued for a new way of looking at the non-epistemic status of deference in a variety of areas. By examining different cases of deference, I have shown that the intuitive asymmetry of deference needs to be specified more carefully. Thinking about problematic and unproblematic deference as mapping onto domains, however narrow or broad, does not do justice to our practices and intuitions regarding when it is unproblematic to take testimony. I have proposed instead that we can distinguish between problematic and unproblematic deference by looking at whether it interferes with our capacity for practical deliberation. If it does, then deference is *pro tanto* bad, because practical deliberation is instrumentally and (extrinsically) finally valuable, and it is *pro tanto* bad to interfere with something of value. I have tested this criterion against a series of examples and it turns out that the cases in which the capacity is interfered with are precisely the cases about which we have negative intuitions. Thus, my practical deliberation view grounds a new version of the asymmetry thesis and as such it is the most developed and broad view on the problematic character of deference that currently exists in the literature.

Chapter VI: Moral Deference and Metaethics

In this chapter I investigate the relation between moral deference and metaethics.¹⁸⁵ In the previous chapters of this thesis I have engaged with different pessimist arguments that wished to explain the problematic character of deference. Moreover, I have put forward one such view myself. One thing all these accounts have in common is that they aim to offer an explanation of the problematic character of moral and certain other types of deference by reference to certain (putative) facts about what is valuable, such as moral understanding, authenticity or, indeed, practical deliberation. We can call these *first-order accounts*, since they roughly give recommendations as to what is good or bad when it comes to moral deference, by reference to some values that we hold (i.e. moral deference is good or bad in a certain respect, in certain circumstances).

However, some philosophers have pointed out that we might want to consider moral deference outside the first-order level discourse. That is to say, we can think about moral deference from a *metaethical* perspective because there are some (at least *prima facie*) interesting connections between the two. On the one hand, the problematic character of moral deference seems to pose a problem for certain metaethical views. More precisely, it can be represented as an objection to moral realism. On the other hand, it seems that some metaethical views have the tools to (*prima facie*) easily explain our attitudes towards moral deference. Call these *second-order accounts*, since they do not appeal to what is valuable, but rather look into the foundations of morality to find an explanation for the oddness of moral deference. In sum, there is a connection between how we see moral deference and metaethics: our attitudes towards moral deference might show us something about what metaethical commitments we should hold and our metaethical commitments might help us explain our attitudes towards moral deference.

I will argue that this approach is misguided and that we need not be concerned with metaethics when it comes to moral deference. The problem with moral deference does not lie on a second-order, metaethical level, but rather on

¹⁸⁵ My target is always *recurrent* moral deference, but I do not think that whether or not we defer repetitively makes a difference from a metaethical point of view. So I will omit the ‘recurrent’ qualification for this chapter.

a first-level, axiological level.¹⁸⁶ Our explanation of its suspicious character should be compatible with our most plausible metaethical views. I end by showing how my practical deliberation argument does exactly that: it holds robustly across a range of metaethical views.¹⁸⁷

1. The Oddness of Moral Deference as an Objection to Moral Realism

Why should we even think about drawing a connection between moral deference and metaethics? One reason is that our negative attitudes towards moral deference seem to present a problem for moral realism. Moral realism is the view that there are moral facts that exist independently of our attitudes and so are objective. Our moral judgements are truth-apt, representational beliefs, which aim to capture reality. In this framework, there seems to be no reason why some people could not be in a better position to know these objective moral facts, and thereby come to be in a position to transmit such knowledge to people who lack it. That would mean that some people are moral experts relative to other people. Though nothing about moral realism strictly entails this, it would seem like a likely scenario if moral realism were true.

In *this* sense then, under moral realism, morality is similar to physics, history, chemistry, medicine etc.: it is an area about which we can learn different facts that are out there, from experts.¹⁸⁸ However, deference about, e.g., facts of physics generally seems perfectly fine, whereas deference about moral facts does

¹⁸⁶ I make here the assumption that there are first-order and second-order moral claims, but I do not take a stance on the relationship between them; they could be interrelated or independent from each other. This does not matter for my purposes for reasons that will become clear later.

¹⁸⁷ Discussions about moral deference and metaethics in the literature vary from less to more developed. Hopkins (2007) and McGrath (2009; 2011) provide good starting points: they mention some potential metaethical explanations for our attitudes towards moral deference but, unfortunately, they lack nuance. Fletcher (2016), however, does take this issue seriously. He develops a non-cognitivist account to explain the fishiness of moral deference, as we will shortly see. McGrath (2009; 2011) and Blanchard (2019) also take the possible challenge that the inappropriateness of moral deference presents for moral realism seriously. I discuss that later too.

¹⁸⁸ Note that this does not exclude the non-naturalist. Even if moral facts are significantly different from non-moral facts (e.g. physics facts), different people can still have different levels of moral knowledge about the existing moral facts, and so we can have moral experts.

not. We have no negative intuitions about deference on physics related matters,¹⁸⁹ but we have negative intuitions about moral deference. But, because of the abovementioned similarity between the two, a moral realist should say that moral deference is fine as well. This gives the wrong prediction about moral deference, as our intuitions tell us that it is not fine.

Blanchard (2019, p. 1143) usefully constructs this argument (originally from McGrath, 2011) as follows:

1. If moral realism is true, then there is nothing inappropriate about moral deference *per se*.
 2. There is something inappropriate about moral deference *per se*.
- Therefore,
3. Moral realism is false.¹⁹⁰

The problem is that there is a tension between our intuitions about moral deference and the theoretical implications of moral realism for moral deference. Of course, one way to resolve this tension is to say that our intuitions are wrong and there is actually nothing wrong with moral deference. We have seen in chapter III how the moral deference optimists try to do precisely this: explain away our negative attitudes towards moral deference. We have also seen in chapter III why they are not successful. They do not take into account cases of moral deference that are problematic *per se* and so their explanations are underinclusive. Moreover, in chapter IV, I have given an argument for why our intuitions about moral deference are indeed warranted. Thus, resolving the tension between our attitudes about moral deference and moral realism by dismissing our intuitions will not work. So the immediate conclusion seems to be that moral realism is false (or that this tension is evidence against moral realism).¹⁹¹

¹⁸⁹ This is an oversimplification for the sake of brevity. As I have argued in the previous chapter, things are not always so straightforward when it comes to non-moral deference, and the asymmetry does not cut as neatly. But this will do for the purposes of this chapter.

¹⁹⁰ Blanchard (2019, p. 1144) says that this challenge could also be formulated against cognitivism and objectivism, as constitutive theses of realism. But not every cognitivist view will face this objection, as McGrath (2009, p. 324) points out: a simple subjectivist or certain kinds of error theorists may have the resources to explain our negative attitudes towards moral deference.

¹⁹¹ This objection need not be limited to moral realism, at least at first sight. Some other metaethical theories which allow for moral truths to exist might also have to deal with it. For example, it may challenge some versions of agent relativism or group relativism (but it does not challenge individual appraiser subjectivism or relativism, for reasons given later, even though

Blanchard offers what I take to be a good response to this objection. He argues that the objection equivocates on the meaning of ‘inappropriate’ (Blanchard, 2019, p. 1146). In the first premise, it means *epistemically* inappropriate, whereas in the second premise it means *non-epistemically* inappropriate.

Moral realism commits us to thinking that there are objective, mind-independent facts about which we can form truth-apt beliefs that can be expressed through descriptive assertions. It also suggests that we should hold that some people can be more reliable than others with respect to their moral judgements. These features of moral realism suggest that, according to this view, moral deference should be no more problematic than non-moral deference. But these aspects matter *epistemically* (and perhaps metaphysically and semantically in the big picture) and, importantly, do not have much to say about the non-epistemic status of normative practices, such as moral deference (Blanchard, 2019, p. 1147). A moral realist may well say that, *epistemically*, moral deference is the same as non-moral deference. But this leaves open the possibility that there is something that is non-epistemically wrong with it.¹⁹²

This is precisely how we should interpret premise 2, according to Blanchard: it refers to the non-epistemic inappropriateness of moral deference. He gives three reasons for this. First, there does not seem to be any insurmountable obstacle in being able to transmit moral knowledge through testimony. Testimony, in general, is a perfectly good medium to pass on

moral truths exist on those views). Depending on how they spell their minimal conception of truth, perhaps even some variants of quasi-realism might have to confront this objection.

¹⁹² McGrath’s (2009) strategy of responding to the objection on behalf of the realist goes roughly along these lines (following Anscombe). That moral testimony works epistemically in the realist framework does not mean that it does not face other problems that could explain our intuitions. She argues that knowledge-how is important in morality, as we need it to know *how* to apply moral principles in different situations and how to correctly interpret what the experts are saying such that we can do the right thing in various circumstances. But moral deference can only give us knowledge-that. Knowledge-that does not give us the moral sensitivity that we need to be able to recognize how principles apply in a variety of situations. When it comes to deferring about what to do in particular cases rather than about more general principles, McGrath broadly follows the moral understanding view: acting on the basis of moral deference makes our actions less valuable because the agent does not do the right thing for the right reasons. This means that someone who acts based on deference falls short of the ideal of the virtuous person. This is what explains our negative intuitions towards moral deference as well as the asymmetry with non-moral deference: there is no analogous failure of achieving an ideal of agency when it comes to non-moral deference. This way one can be a moral realist and conserve a negative attitude towards moral deference.

knowledge and *prima facie*¹⁹³ there is no reason to think this would not work — at least in principle — when it comes to moral matters. A person can gain moral knowledge through testimony and, if they do defer, they have behaved in an epistemically appropriate way: no bad blood here, epistemically speaking. Secondly, we have some intuitive data. Blanchard (2019, p. 1148) argues that if the problem with moral deference was in fact epistemic our intuitions would plausibly be different:

The cases [of moral deference] do not resemble, say, Gettier cases, in which there is a widespread intuition that some specific epistemic good, for example, knowledge, is lacking. In those cases, the presence of epistemic luck directly drives the intuition that there is something distinctly epistemically amiss, with (as far as I know) no one thinking that there is anything non-epistemically (say, morally) amiss in these cases. But in cases of moral deference, there is no such distinctly epistemic feature.

Our intuitions do not seem to track any epistemic issues. Instead, they appear to target the agents who defer and their behaviour, and do not rely on any distinctly epistemic features of the case. And, finally, the most promising contenders for explaining the inappropriateness of moral deference that are available are non-epistemic views. There is the moral understanding strategy, the authenticity arguments, and so on. Blanchard holds that as long as one of the existent non-epistemic accounts is true — which seems more likely than the possibility of an epistemic explanation to be true — then, of course, the issue with moral deference will not be epistemic, and so the objection against moral realism will fail.

All these considerations, Blanchard argues, suggest that when we say that moral deference is problematic or inappropriate, we mean non-epistemically problematic or inappropriate. But moral realism only predicts that moral deference does not have to be epistemically problematic or inappropriate. So, in the end, there is no tension between the normative status of moral deference and the predictions of moral realism. I take this to be a good reply to the objection

¹⁹³ I say *prima facie* because, e.g., a convinced non-cognitivist would probably not share this intuition. But see the following point for why even the non-cognitivist cannot sustain the weight of this claim on their shoulders.

against moral realism such that I do not have anything to add about it here. Instead, I will take it further. That is, I will later build on Blanchard's view to offer a general argument against any claim that a metaethical theory purports to make about moral deference. But until then let us see what other metaethical views have to say about moral deference.

2. Metaethical Explanations for the Oddness of Moral Deference

In contrast to moral realism, other metaethical views not only are not challenged by our intuitions regarding moral deference, but actually purport to be able to explain them. Call them *second-order explanations* of the oddness of moral deference, i.e. metaethical explanations, which rely on some aspects related to the nature of moral facts, judgements, or discourse. I consider the main theories that might offer such explanations: *individual appraiser subjectivism*, *non-cognitivism*, and *eliminativist error theory*. I will offer broad sketches of how these arguments could go to show why we have reason to think, at least *prima facie*, that there *could* be such metaethical explanations of the problematic character of moral deference.

I begin with non-cognitivism, the view that moral judgements are some broadly desire-like states, such as preferences, sentiments like approval and disapproval, or emotions. Non-cognitivism usually goes together with expressivism, and so moral statements are considered to not represent or describe the world, but instead to express some non-cognitive states. Typically, this entails the claim that moral statements are not truth-apt (although not always: some varieties have a minimalist conception of truth, e.g. quasi-realism). Philosophers have already briefly noted that at least some versions of non-cognitivism may provide a pessimist explanation for the problem of moral deference:

In sincerely uttering the sentence, "Eating meat is immoral," you are expressing your own negative emotional response toward eating meat (Ayer 1936). Alternatively, perhaps you are expressing your "unfavorable interest" in eating meat (Stevenson 1944), or your acceptance of norms that prohibit the practice (Gibbard 1990). But in any case, one thing that you are

not doing is attempting to depict some moral fact or state of affairs. (McGrath, 2009, p. 321)

The idea is that since moral statements are not the sort of thing that can be true or false, then how could they be transmitted through testimony? What you do when you defer to someone's testimony is taking knowledge from them. But when I say that eating meat is immoral, all I am doing is, e.g., expressing my own emotional reaction to eating meat, and so there are no moral truths or knowledge to be passed on.¹⁹⁴ But, as others have noticed (Hills, 2013, p. 559, endnote 4; Fletcher, 2016, p. 59), this explanation is available only to those non-cognitivists that do not believe that there is such a thing as moral knowledge. Denying that there is moral knowledge is, of course, controversial and can be problematic. Such views, for example, cannot account for genuine moral disagreement in belief (Sliwa, 2012, p. 176).

But more sophisticated versions of non-cognitivism, which allow for moral knowledge, might also be able to offer an explanation. Fletcher (2016) develops such a view. We know that direct moral deference about pure assertions implies that we supposedly adopt a judgement about purely moral matters solely on the basis of another person's moral testimony. Since moral judgements in a non-cognitivist framework are desire-like attitudes or sentiments, that would mean that when we defer we would somehow have to adopt as our own the desire-like attitudes or sentiments of someone else. But, "moral sentiments are at least difficult to form on the basis of pure, direct, testimony" (Fletcher, 2016, p. 60). Fletcher gives three reasons for this claim.

One, moral sentiments are formed in response to certain features of an object and one needs to at least be aware of those features to feel the appropriate sentiments. Direct testimony about pure assertions does not provide such details. So, even if I know I should have anger towards John, I will not be able to feel angry because I lack the experience of the features that would prompt anger. Secondly, moral sentiments are formed in response to certain features of an object and, as such, we take them to be merited by the person or situation and, in this sense, authoritative (if there is no evidence against that). If John does something wrong, I become angry at him, in response to his behaviour (which I

¹⁹⁴ Hopkins (2007, pp. 615-618) also considers this explanation. He rejects it by saying that it does not do justice to the importance of moral reasoning.

take as evidence of the wrongness of his behaviour, if nothing contradicts that). But if I defer to you about the wrongness of John's behaviour I will not be able to be angry at him because I cannot experience such a reaction as merited — since I do not have access to those features of his behaviour but I only have some testimony. Thirdly, it is not easy at all to change our moral sentiments by will, simply by acquiring some information from testimony (which contrasts with beliefs in (at least some) cases where I acquire new information). I cannot suddenly get angry simply because I want to get angry at John, as someone told me I should.¹⁹⁵

So in a non-cognitivist framework this is why moral deference seems problematic: moral judgements cannot be formed on the basis of pure, direct deference. They are the kind of thing that is at least difficult to adopt through deference. Fletcher (2016, p. 63) thinks this is analogous with being told by a food critic that a certain food that you find disgusting is actually delicious. You cannot change the way you feel about the food only because you have been told you should feel something else. In the case of moral judgements, “it is difficult to feel anger or resentment towards capital punishment simply because someone (even someone you trust) tells you that such responses are correct” (Fletcher, 2016, p. 63). His account is then based on the *psychology* of forming moral judgements. In particular, the point is that pure moral deference is at least very difficult to occur (if not impossible). This is also why it is so rare, according to him (2016, p. 54).

Non-cognitivism is not the only metaethical view that might explain our attitude towards moral deference. Individual appraiser subjectivism is another candidate. On this view, judgements about what is morally right depend on what the appraiser desires or approves of (Van Roojen, 2015, p. 99). Moral judgements are cognitive states, representational (both semantically and psychologically) and

¹⁹⁵ Fletcher (2016, pp. 65-67) holds that his argument can apply under hybrid metaethical theories as well, i.e. any view that has desire-like moral sentiments as at least partly constitutive of moral judgements. Moreover, some forms of moral realism can also work with his view. For example, on a realist-sentimentalism view, moral judgements are beliefs about the fittingness of moral sentiments. This compels one to have their beliefs align with their sentiments. If they do not, there will be some disharmony or even rational incoherence within the agent, which will beg to be solved in some way. Since moral deference will likely not help one have their beliefs match up with the fitting sentiments, moral deference will be problematic.

truth-apt beliefs, but they do not track any objective¹⁹⁶ mind-independent moral facts that are out there in the world.¹⁹⁷ Instead they refer to the speaker's subjective mental states and their responses to different situations. At the level of statements, 'Stealing is wrong' implicitly bears the qualification 'for me', as it essentially refers to what *I* think of stealing, i.e. whether I approve or disapprove of it. The utterance is true or false depending on whether I approve or disapprove of stealing.

If these are the commitments of the view, then it is not difficult to see why moral deference will be problematic in this framework. Moral truths are subjective and depend on what each of us approve or disapprove of, which might be different from person to person. If you just take my testimony, you are only learning what *I* approve of. If for a subjectivist 'Stealing is wrong' needs to be qualified with 'for me', then deference will confuse the indexical by mistakenly attributing that belief to you when it actually belongs to me, the person who genuinely gives the approval. Moral statements will always be relative to the speaker on this metaethical view, which renders moral beliefs adopted on deference at best false and at worst almost senseless. That is to say, false because they express someone else's approval or disapproval. It is not true that *you* disapprove of stealing, if you just deferred to me. *I* uttered the statement, because *I* disapprove of stealing. Alternatively, beliefs formed on the basis of deference are almost senseless because they would miss their essential component, i.e. the speaker's own approval or disapproval. Is it even possible to make a moral judgement when *my* approval or disapproval is not actually present? So, in brief, the problem with moral deference on an individual appraiser subjectivist view is that the subjective nature of moral properties and facts makes it difficult to pass on moral knowledge through testimony. McGrath (2009, p. 323) interprets it as

¹⁹⁶ This might depend on how we interpret this objectivity. Even though moral judgements on this view are not universal, "the fact that some person's or group of people's values forbid the action would still seem to be a fact that is fully out there and not depending metaphysically on the "projection" or "construction" of human thought" (Chrisman, 2017, p. 72; although this fact might be a psychological one). I will, however, discuss objectivity here in the way robust moral realists do: as having the property of universality and hold that objective moral judgements will be the same for everyone, from any perspective.

¹⁹⁷ A subjectivist need not necessarily be a cognitivist and a representationalist, though the simple subjectivist I discuss here is usually characterised as such (Van Roojen, 2015, p. 117). Instead of holding that moral sentences are about approval or disapproval, they could say they *express* those states. Such versions of subjectivism would then just fall under the general non-cognitivist umbrella.

“a special case of the oddity of deferring to someone else about one’s own mental states”. That is because each speaker is the most qualified person to say whether they approve or disapprove of something rather than anyone else.

Finally, there is the eliminativist moral error theory which also has a way of making us wary of moral deference.¹⁹⁸ According to this view, our moral discourse aims to describe or represent reality, but ultimately fails because reality just is not the way it is portrayed by it. Error theorists say that all our moral judgements are untrue insofar as they purport to be about objective and irreducibly normative moral facts, as such facts do not exist. When we talk about morality, we do it by presupposing the existence of some kind of facts or properties that exist in the fabric of the world and which are not dependent on us. But such things do not actually exist. These two claims — the semantic and the metaphysical-ontological one — point to the conclusion that our moral discourse is untrue. And eliminativists want to eliminate such discourse: they hold that we should stop talking about morality the way we do.

Moral deference is thus problematic because it is a practice that keeps us persisting in error, by enabling the trade of false beliefs between us. An eliminativist error theorist would probably say that we should not even *consider* the practice of deference: what we should do is to stop participating and instead get rid of our current moral discourse altogether. The problem seems at least *prima facie* epistemic: it has to do with the badness of transmitting false beliefs, insofar as it is better to have true rather than false beliefs, perhaps both epistemically and pragmatically.¹⁹⁹ But it is not just this; I think the worry goes deeper. The error theorist does not just warn us of the possibility of gaining some false beliefs when it comes to morality. Rather they want to denounce the moral discourse as a whole and uncover that it is based on a systematic error, and moral deference goes against that, by contributing to its persistence.²⁰⁰ In a similar vein, on some versions of moral scepticism one could also say that, since we do not

¹⁹⁸ Error theorists can also choose to be, e.g., fictionalists and think that we should keep our current moral discourse because it is a useful act of pretence. Here I set aside non-eliminativist versions of error theory because their capacity to give a metaethical explanation of why moral deference is odd is much less clear.

¹⁹⁹ It is definitely not the *moral* badness of spreading false beliefs that is relevant here because that would obviously not hold in an error theory framework.

²⁰⁰ McGrath (2009, p. 330) takes error theory to be interchangeable with nihilism, and so claims that the same explanation holds for the nihilist.

have moral knowledge (or justification), there is nothing to be transmitted through testimony, and that is why moral deference is problematic (McGrath 2009, p. 329).²⁰¹

Now, to take stock: we have just seen, in broad brushstrokes, some ways in which various metaethical views could provide explanations for the problematic character of moral deference. And here is why this matters. If it is possible to give a metaethical explanation to the problematic character of moral deference, then at least at first sight it seems that first-order accounts, such as my practical deliberation argument, have a potentially strong rival. We have a second-order explanation for our intuitions, so there is no need to study the problem at the first-order level. Additionally, the fact that our attitudes towards moral deference are purported to put pressure on our metaethical commitments (like the objection against moral realism) also suggests that the problem belongs to the second-order metaethical level. However, I do not think this is the case. The metaethical stories related to moral deference are at best incomplete and at worst irrelevant.

3. Metaethical Explanations for the Oddness of Moral Deference: A Response

In what follows I construct an argument inspired by Blanchard's (2019) response to the objection to moral realism to show the possible metaethical explanations discussed previously do not tell the whole story about moral deference. As such, they do not present a threat to my practical deliberation view.

The argument is the following:

²⁰¹ McGrath (2011, p. 117) also considers a constructivist explanation, according to which the problem is that someone who defers seems to implicitly subscribe to the view that there are mind-independent facts that our moral judgements try to represent. However, this is a misleading picture of morality, she says, since on this view moral facts are something to be constructed rather than discovered.

P1: Metaethical views do not, as such, explain first-order facts.

P2: Our negative intuitions about moral deference suggest that there is a first-order problem with it.

C: Metaethical views cannot explain our negative intuitions about moral deference.

The main idea is this: both the purported metaethical explanations for our intuitions about moral deference and the objection to moral realism operate on a second-order level, mainly by appeal to aspects related to the possibility of moral knowledge, as well as to the nature and characteristics of moral facts and of our moral judgements. They give us no insight into the appropriateness of our first-level practices. But our intuitions about (and the real problem with) moral deference are not about second-order issues. Or, even if we would accept that some of them are, they do not exhaust the range of relevant intuitions. The fishiness of moral deference persists robustly regardless of what our metaethical commitments are (if we have any). This suggests that the philosophically fundamental issue lies somewhere else: at an axiological first-level, being connected to what is valuable for us as humans. So, any relation that we attempt to draw between our attitudes towards moral deference and metaethics misses the point. In this sense, my argument is inspired by Blanchard's: it shows that there is a different focus in the metaethical explanations that are allegedly available and our intuitions regarding moral deference. Now allow me explain why.

I start by justifying P1. The first thing to say is that, although it is possible that there could be connections between our normative views and our metaethical theories, it is a mark of the latter that it is concerned with “the metaphysical, epistemological, semantic, and psychological, presuppositions and commitments of moral thought, talk, and practice” (Sayre-McCord, 2014). That is, metaethical claims are *second-order* claims about the fundamentals of morality, moral thought, talk, and practice. They are not (at least for the most part) concerned with our first-order ethical business, so to speak. The job of metaethics is not to tell us what action or behaviour or practice is right or wrong; what is valuable and what values we should pursue; whether agents are morally

defective or praiseworthy in any way. This is what our first-order normative theory does. So any metaethical explanation of moral deference will operate on a second-order level simply because that is the main jurisdiction of our metaethical theories. In this sense, it will not be able to explain the way in which moral deference is non-epistemically, normatively or evaluatively, problematic that we are concerned with, as I will show later.

I do not want to reiterate section 2 from above, but recall the particulars of the attempts made by different metaethical theories to explain our negative intuitions about moral deference. According to individual appraiser subjectivism, moral judgements are about the individual's states of approval or disapproval. So, because of the kind of things moral judgements are, forming a belief based on another person's judgements, i.e. another person's approval or disapproval, will be at least epistemically fishy in some way. Such a belief would be false (we lack the actual approval or disapproval), or at least epistemically inferior, if we think this is like deference about self-knowledge (knowing whether you approve or disapprove of something yourself), where no one else can be, *prima facie*, an expert relative to the speaker. An eliminativist error theorist focuses on our moral discourse, with appeal to the existence and nature of moral facts. Our moral discourse, insofar as it purports to represent some objective moral facts or properties — and thus deference as a part of it — is simply untrue and misguided.

Certain (though not many nowadays) versions of non-cognitivism state that the utterances of our moral judgements are expressions of desire-like states, and as such have no truth value. This suggests that moral knowledge does not exist. In such conditions, how could we defer to one another? So if someone is interested in finding out some moral truth or get knowledge — as most, or all, people who seek to defer do — they will not be successful.

Other versions of non-cognitivism (or hybrid theories) claim that as long as there is at least some desire-like state that is constitutive of moral judgements then we will find it at least difficult (or maybe even impossible) to form such a judgement on the basis of testimony (e.g. Fletcher's view). This basically says that there can be no uptake of moral testimony.²⁰² Fundamentally, this seems to be a claim about the nature and the psychology of moral judgements. Specifically, the

²⁰² This, however, does not account for the badness of active deference, i.e. acting on the basis of deference. But then the non-cognitivist explanation would be underinclusive.

idea is that pure moral deference is at least very difficult to actually happen (maybe even impossible).

But, again, these explanations are second-order. Since they are based on metaphysical, semantic, psychological, and epistemological claims, they do not, as such, have anything, or much, to say about our first-order practices. However, this is not in line with our intuitions about moral deference and our best explanations of its problematic character. Rather, they suggest that the problem is not on a second-order level. I give you five reasons why P2 in the argument above — our negative intuitions about moral deference suggest that there is a first-order problem with it — is justified.

Firstly, in line with what Blanchard says, our intuitions do not seem to concern the epistemology of moral testimony. For example, the oddness of moral deference does not go away even when we assume that everything goes well epistemically, i.e. when we imagine a case where moral knowledge exists, it can be transmitted through testimony, and our testifier is knowledgeable, reliable, trustworthy, and so on. Moreover, the possibility of giving a *conditional* non-epistemic explanation of the problematic character of moral deference is coherent as well. That is, it does not seem misguided to ask: *if* everything is epistemically okay with moral deference, are there any non-epistemic reasons not to defer? Or, *if* everything is epistemically okay with moral deference, could there be a non-epistemic explanation for our negative attitudes towards it? So if any metaethical explanation is based on some second-order claims that want to deem moral deference problematic because of some epistemic problems it might have (e.g. there is no moral knowledge), then it would not be a good explanation. At worst, the questions just mentioned would have to be deemed non-sense, which would seem strange. Or, at best, this would show that a metaethical explanation is incomplete and that a first-order non-epistemic account of our intuitions about moral deference would still be welcome.

A similar point holds for psychological explanations of the problematic character of moral deference. A non-cognitivist view like Fletcher's (2016) states that the issue is not that there is no moral knowledge, but that it is difficult to form moral judgements based on deference, since moral judgements are at least partly constituted by moral sentiments. Moral deference rarely occurs, if ever, precisely because of this. But our intuitions do not seem to align with this view,

i.e. they do not appear to track the psychological aspects of moral deference. It is implausible that our intuitions about moral deference are to be interpreted as reactions of disbelief at hearing that someone deferred because we think that it is at least very difficult, psychologically, to form beliefs based on testimony.²⁰³ Rather, our attitudes seem to be about how it is bad, in some way, that the deferrer did that, or that they should not have done it. We simply assume that we are able to defer and that it is not particularly difficult to do so. Thus, if we *believe* moral deference is possible (even if it is *in fact* not possible), then we still need to explain our intuitions about it in some way or another. We cannot explain them, however, by saying that moral deference is difficult or impossible, precisely because that purported fact is opaque to those of us who do believe that it is possible and that it can occur with relative ease.

Secondly, for one of the metaethical explanations discussed in the previous section to apply, it would mean that all those who find moral deference problematic would have to be non-cognitivists, appraiser subjectivists etc. McGrath (2009, p. 330) makes this point against the moral sceptic, the moral nihilist, and the error theorist who purport to explain the badness of moral deference via their metaethical commitments. If most of us are *not* sceptics, nihilists, or error theorists, we still need to explain why we have a negative attitude towards moral deference, even if nihilism or error theory is actually true. For example, McGrath (2009, p. 330, my italics) compares this with why deference about the results of future coin tosses and lotteries (for which we do not have evidence) is problematic:

the reason it seems bizarre to defer about the outcomes of future coin-tosses and lotteries is not simply the fact that no one is currently in a position to know the outcomes, *but rather that this fact is widely known.*

²⁰³ If this were true, it would also make a lot of us — those who think deference is possible or even fairly easy to happen — just confused. For example, when we make the distinction between cases of moral deference that are problematic and those which are not, or those which are problematic *per se* and those which are not, we are just talking non-sense. Those who think they have deferred on occasion are probably even more confused — they might think they formed moral judgements, when they actually did not (because they are likely to lack the fitting sentiments). It is, of course, possible that all these people are confused. But I do not think they are. There are reasons to believe that moral deference does actually occur. Nevertheless, even if all these people were confused, the task to explain why this *apparent* moral deference seems problematic remains. The debate on moral deference has traction because of our attitudes towards it.

We can imagine that it is *not* widely known that the evidence for these results is inaccessible to us. It is likely that deference about the outcomes will not seem just as problematic as they do now.

I think this explanation extends beyond the views McGrath mentioned to all those I have discussed in the previous section. People would have to believe in the truth of such a view in order to have their negative intuitions towards moral deference explained via that view (or the truth of such a view would have to be widely known). For example, if I do not even know about the existence of a view such as individual appraiser subjectivism, how could my *attitudes* towards moral deference be best explained by something of which I am not even aware? As McGrath (2009, p. 330) puts it: “what matters in explaining our attitude toward pure moral deference is what we believe about the moral facts”.

Moreover, if one of the metaethical explanations mentioned would actually work, that would also mean that moral deference should not trigger negative intuitions for, e.g., a moral realist or a moral fictionalist. Supporters of such views would have to say that moral deference is not only epistemically fine, but fine from any point of view. But this is not the case. Moral deference appears equally problematic in a realist or fictionalist framework, just as it does under non-cognitivism, subjectivism, or eliminativist error theory. This suggests that our intuitions do not hinge upon our metaethical commitments.²⁰⁴

Thirdly, I can lend further support to Blanchard’s claim that our intuitions about moral deference generally seem to target the agents who defer and their behaviour. It seems true that when we see someone deferring we think about what this does to them as moral agents or to their actions, and so that our attitudes are on a first-order level. To see this, think also about how we are sometimes happy with deference. Typically, we are (all-things-considered) happy with deference when we investigate the circumstances of *the agent* and realise they

²⁰⁴ One might say that people’s actual metaethical commitments are not important. The problematic character of moral deference is explained, if one of those metaethical theories is correct, and this is all that matters. However, my argument here is about our *intuitions* and *attitudes* towards moral deference. Our intuitions and attitudes do depend, at least partly, on our commitments, on what we believe, and of what we are aware. One can nevertheless stick to their metaethical guns. But even then we would still need to explain why we have such intuitions, if they do not themselves depend on our metaethical commitments or on the correct metaethical theory itself. This shows, I think, that a metaethical explanation of the fishiness of moral deference would be at least incomplete.

could not have found the answer to their moral problem on their own or that they were impaired in some way and could not decide etc.²⁰⁵

To this, a kind of non-cognitivist, for example, could say that there might be cases where someone does manage to form the relevant kind of mental state through deference. In those cases, no oddity persists and we can be happy with moral deference then. However, these would be (psychologically) anomalous cases. But such an explanation would not allow us to ask whether there is still something *pro tanto* bad about such a case, where it is all-things-considered unproblematic to defer. It would deem this question meaningless or, at least, easily closed with a negative answer. Nevertheless, the question does not seem meaningless and it is possible to give plausible reasons for why some cases of moral deference could be *pro tanto* bad even if all-things-considered good, as I have shown in chapter IV. We have seen throughout this thesis that we can give different *pro tanto* and all-things-considered verdicts in various cases of deference, and diagnosing different instances is not a black and white matter. This, in turn, suggests that the problematic character needs to be discussed on a first-order level. Again, at best, a metaethical explanation is only incomplete.

Fourthly, we seem to have a problem with people who defer about what to do as well as with those who defer about what to believe. However, metaethical explanations can cover only the latter, as they usually do not have much to say about the former. For example, subjectivism and non-cognitivism tell us about the nature of our moral judgements and imply certain commitments about our ethical language. Moral realism and error theory bind us to specific views about our moral discourse and the metaphysical status of moral facts and properties. These sorts of aspects are supposed to explain the fishiness of moral deference. But when I defer about what to do in particular situations — sometimes without even forming a moral judgement at all — how do these aspects even come into play? It is not obvious that actions (which are unaccompanied by judgments) can be metaethically criticised in the same way that moral judgements can.

Finally, I would hope that I have successfully argued that my own practical deliberation argument is able to explain, at least somewhat convincingly,

²⁰⁵ Or when we make the distinction between moral deference that is problematic *per se* or for independent reasons (e.g. character flaw). Metaethical explanations could not cope with this distinction.

the problem of moral deference. As such, most of this thesis is aimed at arguing for the claim that there is a first-order, non-epistemic, with moral deference. And it is a first-order, non-epistemic problem because it is about how the practice of deference can interfere with something which is valuable for us as humans: the exercise and the development of our capacity for practical deliberation. That is, in turn, what makes moral deference *pro tanto* bad. If my argument is to any measure persuasive, but the issue was instead on a second-order, metaethical, then that would be at least a bit odd. It would be a pretty big coincidence that a first-order argument could explain what is bad about deference without appealing to metaethics, where the problem would actually supposed to be.²⁰⁶

Additionally, as I have argued in chapter V, it is not at all clear that the asymmetry thesis can be neatly drawn between moral and non-moral deference, or even normative and non-normative. I have argued that we have reasons to resist such a claim. If that is true, then there might be some cases of non-moral or non-normative deference that could be problematic *per se* and thus not covered by any relevant metaethical or metanormative explanation.

For these reasons, I think, it is plausible to say that the problem with moral deference is a first-order (and non-epistemic) one. But the metaethical explanations I have discussed operate at a second-order level. The supporters of such views would have to think that the issue is about the nature of our moral judgements, about our moral language and discourse, or about moral metaphysics. The badness of moral deference in such a framework is very different from the badness of moral deference that, e.g., my practical deliberation argument purports. If both me and these metaethicists would talk about how moral deference is problematic, we would be equivocating the word ‘problematic’ as well (just like Blanchard accuses the enemies of realism), because we would mean different things by it. The metaethical explanations I have presented earlier in this chapter do not target the relevant problem of deference, namely the one that our intuitions track.

²⁰⁶ Blanchard (2019, pp. 1148-1151) makes the point that the best explanations for the problematic character of moral deference that exist in the literature are non-epistemic, e.g. the moral understanding strategy; the autonomy explanation; the authenticity argument. As long as one of the existent non-epistemic views is true — which seems more likely than the possibility of an epistemic explanation to be true — then obviously the issue with moral deference will not be epistemic.

4. An Objection

Although what I have said in the previous section might be true, the non-cognitivist could have a comeback. They can explain the first-order, non-epistemic problems of at least some fishy cases of deference. For example, the following case triggers negative intuitions: I blame, resent, or am angry at John for eating meat because you blame him. That seems suspicious. The non-cognitivist can easily explain why: it is at least difficult to form moral sentiments or desire-like attitudes, on the basis of deference. Non-cognitivism would then have all the necessary elements: an issue about moral sentiments that triggers negative intuitions at the first-order, non-epistemic, level, and an explanation based exactly on the nature of those moral sentiments.

One problem with this is that it does not sufficiently generalise. For one, as I have argued in the previous sections, deferring on moral judgements such as ‘eating meat is morally impermissible’ cannot be explained by non-cognitivism at the relevant level. Secondly, this account cannot handle active deference, namely deference about what to do (not about what to believe). Not eating meat (without forming the judgement that eating meat is morally impermissible) because someone told you to do so is at least somewhat problematic. But in the non-cognitivist framework only moral judgements and moral sentiments are difficult to form on the basis of deference. The view cannot prescribe anything similar for actions performed on the basis of deference. Thus the view does not have the resources to explain the problematic character of active moral deference. In this sense, the view is underinclusive and incomplete.

Another problem is that this non-cognitivist explanation is not necessarily the best explanation for our intuitions about the oddity of forming moral *sentiments* on the basis of deference. Lord’s (2018) acquaintance argument would also work perfectly well, for example — even in a realist framework. Recall that on his view moral deference is problematic because it cannot give us appreciative knowledge and the appropriate conative and affective attitudes that come with it. We cannot have the full range of appropriate rational reactions to some normative properties if we form judgements on the basis of deference. For example, we cannot feel anger or resentment of the basis of deference because we lack acquaintance. This view would then explain everything that needs

explaining at the first-order level, without having to commit to a certain metaethical view (but even on a robustly realist view).²⁰⁷ Of course, I hope to have shown that my practical deliberation view is preferable to Lord's view, but the point here is that even if we were inclined to have an explanation that is based on issues that have to do with moral sentiments, there is no need to choose a metaethical one. Better alternatives are available.

5. Concluding remarks

To end, I want to mention that my practical deliberation argument is a first-order, non-epistemic explanation of the problematic character of (moral) deference that by itself carries no particular commitments to any metaethical (or meta-normative) theories. My view can hold robustly across various such views, like the ones I have discussed in this chapter. Practical deliberation can be considered valuable in any metaethical framework.

My view puts no pressure on the realist's metaethical commitments in any way. On the contrary, practical deliberation can be considered important on a realist approach to morality. For one, it could be considered important for the reasons I outline in chapter IV, i.e. its instrumental and extrinsic final value. Further, it could be considered particularly important if we think that we can learn about moral facts through deliberation. Then we would certainly have an incentive to develop it, as it would have some added value that is specifically salient for the realist. Either way, being a moral realist does not entail that one should not take our capacity for practical deliberation to be valuable.

A non-cognitivist and a subjectivist can also accept my view too. Our capacity for practical deliberation can be seen as valuable on any non-cognitivist view. Some objects or situations are not present in front of us or detailed enough; some are complex and might not evoke in us non-cognitive reactions straight away. Perhaps we need to deliberate about what kind of feelings and attitudes we

²⁰⁷ One might say that fundamentally Fletcher's view is also a variation of an acquaintance-based view. In a sense, it is. But it also carries some serious metaethical commitments, i.e. a commitment to non-cognitivism, which makes it implausible in some sense. For example, as discussed above, it suggests that moral deference is very difficult to occur or perhaps even impossible. At least Lord's view is a bit lighter on the metaethics and it would be preferable on parsimony grounds.

have or should have towards them. Moreover, we will always need, at least in some situations, to deliberate about what *to do*. Whatever one thinks about the moral judgements that might go along with an action, we will perform morally charged actions that we might need to deliberate about, because metaethics does not tell us what ought to be done. We have to figure that out ourselves, regardless of our metaethical outlook. Practical deliberation helps us with that, so a non-cognitivist can still deem it valuable. And the same sort of thing goes for individual appraiser subjectivism.

For an eliminativist error theorist this might not be so clear. That is because they would presumably reject any moral judgement as mistaken moral (plus perhaps also any normative judgement, if they are a normative error theorist).²⁰⁸ However, even an error theorist will have to do things and to act, in general, in life — so practical deliberation will be at least useful for them (even if not valuable). So perhaps they could see recurrent moral deference as *pro tanto* bad in the sense of it not being useful. This feels like a stretch, however, and I might have to end by admitting that my view does not really work in a moral or normative eliminativist error theory framework. However, I do not mind it that much. If error theory is true, then we should not even be talking about any of this anyway.

²⁰⁸ But if one is a moral rather than normative error theorist then they might accept that practical deliberation has value (albeit not moral value).

Conclusion

In this thesis I put forward a new version of moral deference pessimism and advanced an argument that aims to explain the non-epistemically problematic character of moral deference *per se*, and deference (*per se*) more generally. I analysed the central concept — that of moral deference — and reconstructed the debate in the clearest way I could in chapter I. In chapter II and III, I discussed the existing moral deference pessimist and optimist views. I presented some challenges to those views in order to motivate my project. I do not pretend I showed those views are plainly wrong. In some places, I asked questions that they may well answer and in other places I offered suggestions of improvement that they may well take on board. But I do think that my arguments give us at least some reason to think that the discussion on moral deference is not finished, and that there is space for another view that might do better. In chapter IV, I put forward such a view: my practical deliberation argument. In chapter V, I explored the potential of this argument beyond moral deference and showed that it could be ambitious enough to give us a principled way of distinguishing between problematic and unproblematic deference more generally. In chapter VI, I investigated the connections between moral deference and metaethics, and defended the idea that a metaethical explanation of our attitudes towards moral deference is not the way to go.

The view that I advance in this thesis is not a hundred percent complete. I present it as an alternative to the existing (moral) deference pessimist views, fully aware that it can be challenged in many ways (as any argument can). I do not purport to have all the answers. Rather, I am just exploring how far it can go. I hope to have shown that it can go pretty far. As such, I do think that it is a candidate worth having at the table. So, it turns out, my mother was only sort of right, with respect to some cases.

This project, with the arguments and the results it contains, has opened some interesting avenues of research that I take as worthwhile to pursue in the future. Roughly, they are as follows. First, there is the question of whether the practical deliberation account I present here for the problematic character of deference could have an analogue when it comes to theoretical deliberation. That is, whether certain cases of deference could interfere with the exercise and

development of the capacity for theoretical deliberation, such that theoretical deference is also non-epistemically bad. Further investigation could be done regarding the respects in which theoretical and practical deliberation are similar, as well as different, and what implications this has for whether recurrent deference can interfere with both. The trick would be to develop an account that can offer a principled distinction between the cases of deference that are problematic and those which are not, in a way that aligns at least broadly with our intuitions. We should not want to end up telling people that they have reason to refuse most testimony!

Secondly, I believe there are fascinating issues to be investigated concerning some special cases of moral deference, by exploring the relationship between practical deliberation and feminist standpoint theory. Do some of us have first-personal epistemic access to certain matters such that we are epistemic authorities on that and can better deliberate about them? For example, perhaps women have a privileged access to issues that specifically concern them or are lived exclusively by them. I think that how we answer such a question can shed new light on thinking about moral deference involving matters that concern particular groups, e.g. racism, sexism, classism, ableism. If standpoints specifically contribute to practical deliberation, by providing some privileged epistemic access to some people, then we need to reflect on whether this provides reasons — and, if yes, what kind and how strong — in support of deference to such people (this builds upon Wiland, 2017 — see chapter III).

Finally, in this thesis, I make a distinction between cases where we could deliberate such that we reach an answer to our conundrums and cases where we cannot. But questions arise: *when* does practical deliberation count as being out of reach? How much should we try and deliberate before deferring? Are we to blame if we do not make enough of an *effort*? I believe that answering these questions could broaden the ongoing debate about blameworthy actions to practical and moral *deliberation*. An interesting path to take would be to investigate whether we might be blameworthy for not trying sufficiently to deliberate on our own, at least at least regarding moral matters. Further research could be done by seeing whether there is anything that would justify such blame and how we can distinguish between instances where someone is blameworthy and where they are not.

Bibliography

- Adler, J. 2017. Epistemological Problems of Testimony. [Online]. In: Zalta, E. N. ed. *The Stanford Encyclopedia of Philosophy*. [Accessed 22 August 2019]. Available from:
<https://plato.stanford.edu/archives/win2017/entries/testimony-episprob/>.
- Aristotle. 2000. *Nicomachean Ethics*. Cambridge: Cambridge University Press.
- Arpaly, N. 2003. *Unprincipled Virtue*. Oxford: Oxford University Press.
- Arpaly, N. and Schroeder, T. 2012. Deliberation and Acting for Reasons. *Philosophical Review*. **121**(2), pp. 209-239.
- Audi, R. 2004. Reasons, Practical Reason, and Practical Reasoning. *Ratio*. **17**(2), pp. 119–149.
- Best, D. 1982. The Aesthetic and the Artistic. *Philosophy*. **57**(221), pp. 357-372.
- Blanchard, J. 2019. Moral Realism and Reliance on Moral Testimony. *Philosophical Studies*. **176**(5): 1141-1153.
- Broome, J. 2013. *Rationality through Reasoning*. Oxford: Wiley Blackwell.
- Burns Coleman, E. 2013. Etiquette: The Aesthetics of Display and Engagement. *Aesthetics*. **23**(1), pp. 68-91.
- Buss, S. 1999. Appearing Respectful: The Moral Significance of Manners. *Ethics*. **109**(4), pp. 795–826.
- Callahan, L. F. 2018. Moral Testimony: A Re-Conceived Understanding Explanation. *Philosophical Quarterly*. **68**(272), pp. 437-459.
- Carroll, N. 2000. Art and Ethical Criticism: An Overview of Recent Directions of Research?. *Ethics* **110**(2), pp. 350–387.
- Charness, N., Tuffiash M., Krampe R. T., Reingold E., and Vasyukova, E. 2005. The Role of Deliberate Practice in Chess Expertise. *Applied Cognitive Psychology*. **19**, pp. 151-165.
- Cholbi, M. 2007. Moral Expertise and the Credentials Problem. *Ethical Theory and Moral Practice*. **10**(4), pp. 323-334.
- Chrisman, M. 2017. *What Is This Thing Called Metaethics?* Oxon: Routledge.
- Christiano, T. 2013. Authority. [Online]. In: Zalta, E. B. ed. *The Stanford Encyclopedia of Philosophy*. [Accessed 22 August 2019]. Available from:
<https://plato.stanford.edu/archives/spr2013/entries/authority/>.

- Covaci, A. 2015. *Is Morality for the Chosen Few? The Problem of Moral Deference*. MA thesis, Central European University.
- Cox, D., La Caze, M., and Levine, M. 2018. Integrity. [Online]. In: Zalta, E. N. ed. *The Stanford Encyclopedia of Philosophy*. [Accessed 22 August 2019]. Available from: <https://plato.stanford.edu/archives/win2018/entries/integrity/>.
- Crisp, R. 2014. II—Moral Testimony Pessimism: A Defence. *Aristotelian Society Supplementary Volume*. **88**(1), pp. 129-143.
- Cross, A. 2017. Art Criticism as Practical Reasoning. *British Journal of Aesthetics*. **57**(3), pp. 299–317.
- Cross, B. 2017. Autonomy and Moral Deference. *South African Journal of Philosophy*. **36** (2), pp. 278-291.
- Driver, J. 2006. Autonomy and the Asymmetry Problem for Moral Expertise. *Philosophical Studies*. **128**(3), pp. 619-644.
- Driver, J. 2013. Moral Expertise: Judgement, Practice, and Analysis. *Social Philosophy and Policy*. **30**(1–2), pp. 280–296.
- Driver, J. 2015. Virtue and Moral Deference. *Ethics and Politics*. **17**(2), pp. 27-40.
- Duckworth, A. L., Kirby, T. A., Tsukayama, E., Berstein, E., and Ericsson, K. A. 2011. Deliberate Practice Spells Success: Why Grittier Competitors Triumph at the National Spelling Bee. *Social Psychological and Personality Science*. **2**(2), pp. 174–181.
- Dular, N. 2017. Moral Testimony under Oppression. *Journal of Social Philosophy*. **48**(2), pp. 212-236.
- Edmundson, W. A. 1993. Rethinking Exclusionary Reasons: A Second Edition of Joseph Raz's Practical Reason and Norms. *Law and Philosophy*. **12**(3), pp. 329-343.
- Ehrenberg, K. 2011. Joseph Raz's Theory of Authority. *Philosophy Compass*. **6**(12), pp. 884-894.
- Enoch, D. 2011. *Taking Morality Seriously: A Defense of Robust Realism*. Oxford: Oxford University Press.
- Enoch, D. 2014. A Defense of Moral Deference. *Journal of Philosophy*. **111**(5), pp. 229-258.

- Ericsson, K. A., Krampe R. T., and Tesch-Römer, C. 1993. The Role of Deliberate Practice in the Acquisition of Expert Performance. *Psychological Review*. **100**(3), pp. 363–406.
- Finlay, S. 2019. Defining Normativity. In: Toh, K., Plunkett, D., and Shapiro, S. eds. *Dimensions of Normativity: New Essays on Metaethics and Jurisprudence*. Oxford: Oxford University Press, pp. 187-220
- Fischer, J. M. and Ravizza, M. 1998. *Responsibility and Control: A Theory of Moral Responsibility*. Cambridge: Cambridge University Press.
- Fletcher, G. 2016. Moral Testimony: Once More with Feeling. In: Shafer-Landau, R. ed. *Oxford Studies in Metaethics: volume 11*. Oxford: Oxford University Press, pp. 45-73.
- Fricker, E. and Cooper, D. E. 1987. The Epistemology of Testimony. *Aristotelian Society Supplementary Volume*. **61**(1), pp. 57-106.
- Fricker, E. 1994. Against Gullibility. In: Chakrabarti, A and Matilal, B. K. eds. *Knowing from Words*. Kluwer Academic Publishers, pp. 125-161.
- Fricker, M. 2007. *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press.
- Friedman, J. 2017. Why Suspend Judging?. *Noûs*. **51**(2), pp. 302-326.
- Fritz, J. 2018. What Pessimism about Moral Deference Means for Disagreement. *Ethical Theory and Moral Practice*. **21**(1), pp. 121-136.
- Gaut, B. 2001. Art and Ethics. In: Gaut, B. N. and Lopes, D. eds. *The Routledge Companion to Aesthetics*. 2nd edition. London: Routledge, pp. 341–352.
- Gorodeisky, K. 2010. A New Look at Kant's View of Aesthetic Testimony. *British Journal of Aesthetics*. **50**(1), pp. 53–70.
- Green, C. R. 2008. Epistemology of Testimony. [Online]. *The Internet Encyclopedia of Philosophy*. [Accessed 22 August 2019]. Available from: <https://www.iep.utm.edu/ep-testi/>.
- Grimm, S. 2012. The Value of Understanding. *Philosophy Compass*. **7**(2), pp. 103-117.
- Groll, D. and Decker, J. 2014. Moral Testimony: One of These Things Is Just Like the Others. *Analytic Philosophy*. **54**(4), pp. 54–74.
- Hambrick, D. Z., Oswald, F. L., Altmann, E. M., Meinz, E. J., Gobet, F., and Campitelli, G. 2014. Deliberate Practice: Is That All It Takes to Become An Expert?. *Intelligence*. **45**(1), pp. 34-45.

- Hills, A. 2009. Moral Testimony and Moral Epistemology. *Ethics*. **120**(1), pp. 94–127.
- Hills, A. 2010. *The Beloved Self: Morality and the Challenge From Egoism*. Oxford: Oxford University Press.
- Hills, A. 2012. *Comment on Karen Jones and François Schroeter*. *Analyse & Kritik*. **34**(2), pp. 231-236.
- Hills, A. 2013. Moral Testimony. *Philosophy Compass*. **8**(6), pp. 552-559.
- Hills, A. 2015. Understanding Why. *Noûs*. **50**(4), pp. 661-699.
- Hills, A. 2017a. Aesthetic Understanding. In: Grimm, S. R. ed. *Making Sense of the World: New Essays on the Philosophy of Understanding*. Oxford: Oxford University Press, pp. 159-177.
- Hills, A. 2017b. *Testimony, Understanding, and Virtue in Aesthetics*. Unpublished, pp. 1-19.
- Hoffmann, M. 2012. How to Identify Moral Experts? An Application of Goldman's Criteria for Expert Identification to the Domain of Morality. *Analyse & Kritik*. **34**(2), pp. 299-314.
- Hopkins, R. 2007. What is Wrong with Moral Testimony?. *Philosophy and Phenomenological Research*. **74**(3), pp. 611–634.
- Hopkins, R. 2011. How to Be a Pessimist about Aesthetic Testimony. *Journal of Philosophy*. **108**(3), pp. 138–157.
- Howell, R. J. 2014. Google Morals, Virtue, and the Asymmetry of Deference. *Noûs*. **48**(3), pp. 389–415.
- Hursthouse, R. 2001. *On Virtue Ethics*. Oxford: Oxford University Press.
- Hursthouse, R. and Pettigrove, G. 2018. Virtue Ethics. [Online] In: Zalta, E. N. ed. *The Stanford Encyclopedia of Philosophy*. [Accessed 22 August 2019]. Available from: <https://plato.stanford.edu/archives/win2018/entries/ethics-virtue/>.
- Isserow, J. 2019. Moral Worth and Doing the Right Thing by Accident. *Australasian Journal of Philosophy*. **97**(2), pp. 251-264.
- Jones, K. 1999. Second-Hand Moral Knowledge. *The Journal of Philosophy*. **96**(2), pp. 55-78.
- Jones, K. and Schroeter, F. 2012. Moral Expertise. *Analyse & Kritik*. **34**(2), pp. 217-230.

- Kahneman, D. 2003. A Perspective on Judgement and Choice: Mapping Bounded Rationality. *American Psychologist*. **58**(9), pp. 697-720.
- Kant, I. 1998. *Groundwork of the Metaphysics of Morals*. Cambridge: Cambridge University Press.
- Kauppinen, A. 2018. Practical Reasoning. In: Star, D. ed. *Oxford Handbook of Reasons and Normativity*. Oxford: Oxford University Press, pp. 395-423.
- Keith, N. and Ericsson, K. A. 2007. A Deliberate Practice Account of Typing Proficiency in Everyday Typists. *Journal of Experimental Psychology Applied*. **13**(3), pp. 135–145.
- Kieran, M. 2006. Art, Morality and Ethics: On the (Im)Moral Character of Art Works and Inter-Relations to Artistic Value. *Philosophy Compass*. **1**(2), pp. 129–143.
- Kind, A. 2019. *The Skill of Imagination*. Learning from Imagination, 23-24 May, Leeds.
- Kolnai, A. 1962. Deliberation Is of Ends. *Proceedings of the Aristotelian Society*. **62**(1), pp. 195–218.
- Kulka, T. 1981. The Artistic and the Aesthetic Value of Art. *British Journal of Aesthetics*. **21**(4), pp. 336-350.
- Kvanvig, J. 2003. *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.
- Lackey, J. 2006. Introduction. In: Lackey, J. and Sosa, E. eds. *The Epistemology of Testimony*. Oxford: Oxford University Press, pp. 1-24.
- Lackey J. and Sosa E. 2006. eds. *The Epistemology of Testimony*. Oxford: Oxford University Press.
- Laetz, B. 2008. A Modest Defense of Aesthetic Testimony. *Journal of Aesthetics and Art Criticism*. **66**(4), pp. 355–363.
- Lillehammer, H. 2014. I — Moral Testimony, Moral Virtue, and the Value of Autonomy. *Aristotelian Society Supplementary Volume*. **88**(1), pp. 111-127.
- Lord, E. 2018. How to Learn about Aesthetics and Morality Through Acquaintance and Testimony. In: Shafer-Landau, R. ed. *Oxford Studies in Metaethics: volume 13*. Oxford: Oxford University Press, pp. 71-98.
- Markovits, J. 2010. Acting for the Right Reasons. *Philosophical Review*. **119**(2), pp. 201-242.

- Martin, J. 1993. A Philosophy of Etiquette. *Proceedings of the American Philosophical Society*. **137**(3), pp. 350-356.
- McGeer, V. 2018. Intelligent Capacities. *Proceedings of the Aristotelian Society*. **118**(3), pp. 347–376.
- McGrath, S. 2009. The Puzzle of Pure Moral Deference. *Philosophical Perspectives*. **23**(1), pp. 321-334.
- McGrath, S. 2011. Skepticism about Moral Expertise as a Puzzle for Moral Realism. *Journal of Philosophy*. **108**(3), pp. 111-137.
- McShane, P. 2015. *In Defense of Dependence on Moral Testimony*. PhD thesis, Georgetown University.
- McShane, P. 2018. The Non-remedial Value of Moral Deference. *Philosophical Studies*. **175** (3), pp. 629-647.
- Meinz, E. J. and Hambrick, D. Z. 2010. Deliberate Practice Is Necessary But Not Sufficient to Explain Individual Differences in Piano Sight-Reading Skill: The Role of Working Memory Capacity. *Psychological Science*. **21**(7), pp. 914–919.
- Meskin, A. 2004. Aesthetic Testimony: What Can We Learn From Others About Beauty and Art? *Philosophy and Phenomenological Research*. **69**(1), pp. 65–91.
- Miller, P. 1975. Who Are the Moral Experts?. *Journal of Moral Education*. **5**(1), pp. 3-12.
- Mogensen, A. L. 2015. Moral Testimony Pessimism and the Uncertain Value of Authenticity. *Philosophy and Phenomenological Research*. **93**(3), pp. 261-284.
- Moxley, J. H., Ericsson, K. A., Scheiner, A., and Tuffiash, M. 2015. The Effects of Experience and Disuse on Crossword Solving. *Applied Cognitive Psychology*. **29**(1), pp. 73–80.
- Nguyen, C. T. 2017. The Uses of Aesthetic Testimony. *British Journal of Aesthetics*. **57**(1), pp. 19–36.
- Nickel, P. 2001. Moral Testimony and its Authority. *Ethical Theory and Moral Practice*. **4**(3), pp. 253–266.
- Olberding, A. 2015. It's Not Them, It's You: A Case Study Concerning the Exclusion of Non-Western Philosophy. *Comparative Philosophy*. **6**(2), pp. 14-36.

- Perry, S. R. 1989. *Second-order Reasons, Uncertainty and Legal Theory*. [Online] Faculty Scholarship, Paper 1354. [Accessed 22 August 2019]. Available from: http://scholarship.law.upenn.edu/faculty_scholarship/1354.
- Plunkett, D. and Shapiro, S. 2017. Law, Morality, and Everything Else: General Jurisprudence as a Branch of Metanormative Inquiry. *Ethics*. **128**(1), pp. 37-68.
- Pritchard, D. 2004. The Epistemology of Testimony. *Philosophical Issues*. **14**(1), pp. 326–348.
- Pritchard, D. 2009. Knowledge, Understanding, and Epistemic Value. Royal Institute of Philosophy Supplement. **64**, pp. 19-43.
- Rabinowicz and Rønnow-Rasmussen, T. 2000. A Distinction in Value: Intrinsic and for its Own Sake. *Proceedings of the Aristotelian Society*. **100**(1), pp. 33-51.
- Rabinowicz, W. and Rønnow-Rasmussen, T. 2003. Tropic of Value. *Philosophy and Phenomenological Research*. **66**(2), pp. 389-403.
- Rabinowicz, W. and Rønnow-Rasmussen, T. 2015. Value Taxonomy. In: Brosch, T. and Sander, D. eds. *Handbook of Value: Perspectives from Economics, Neuroscience, Philosophy, Psychology and Sociology*. Oxford: Oxford University Press, pp. 23-42.
- Railton, P. 2004. How to Engage Reason: The Problem of Regress. In: Wallace R. J., Pettit P., Scheffler S., and Smith, M. eds. *Reason and Value: Themes From the Moral Philosophy of Joseph Raz*. Oxford: Clarendon Press, pp. 176-201.
- Raz, J. 1986. *The Morality of Freedom*. Oxford: Oxford University Press.
- Raz, J. 1999. Explaining Normativity: On Rationality and the Justification of Reason. *Ratio*. **12**(4), pp. 354-379.
- Raz, J. 2006. The Problem of Authority: Revisiting the Service Conception. *Minnesota Law Review*. **90**, pp. 1003-1044.
- Reisner, A. and Van Weelden, J. 2015. Moral Reasons for Moral Beliefs: A Puzzle for Moral Deference Pessimism. *Logos and Episteme*. **6**(4), pp. 429-448.
- Richardson, H. S. 1997. *Practical Reasoning About Final Ends*. Cambridge: Cambridge University Press.

- Ridley, A. 2015. Why Ethics and Aesthetics Are Practically the Same. *The Philosophical Quarterly*. **66**(262) pp. 102–121.
- Robson, J. 2012. Aesthetic Testimony. *Philosophy Compass*. **7**(1), pp. 1–10.
- Robson, J. 2018. Aesthetic Testimony and The Test of Time. *Philosophy and Phenomenological Research*. **96**(3), pp. 729–748.
- Rodogno, R. 2015. Prudential Value or Well-being. In Sander, D. and Brosch, T. eds. *Oxford Handbook of Value: The Affective Sciences of Values and Valuation*. Oxford: Oxford University Press, pp. 288–313.
- Rønnow-Rasmussen, T. 2002. Instrumental Values: Strong and Weak. *Ethical Theory and Moral Practice*. **5**(1), pp. 23–43.
- Russell, D. C. 2014. From Personality to Character to Virtue. In: Alfano, M. *Current Controversies in Virtue Theory*. New York: Routledge, pp. 92–105.
- Shapiro, S. J. 2002. *Authority*. [Online]. Stanford/Yale Jr. Faculty Forum Research Paper 00-05, Cardozo Law School, Public Law Research Paper No. 24. [Accessed 22 August 2019]. Available from: <https://ssrn.com/abstract=233830>.
- Sayre-McCord, G. 2014. Metaethics. [Online]. In: Zalta, E. N. ed. *The Stanford Encyclopedia of Philosophy*. [Accessed 22 August 2019]. Available from: <https://plato.stanford.edu/archives/sum2014/entries/metaethics/>.
- Simion, M. 2018. The Explanation Proffering Norm of Moral Assertion. *Ethical Theory and Moral Practice*. **21**(3), pp. 477–488.
- Singer, P. 1972. Moral Experts. *Analysis*. **32**(4), pp. 115–117.
- Skarsaune, K. O. 2016. Moral Deference and Authentic Interaction. *Journal of Philosophy*. **113**(7), pp. 346–357.
- Sliwa, P. 2012. In Defense of Moral Testimony. *Philosophical Studies*. **158**(2), pp. 175–195.
- Sliwa, P. 2015. Understanding and Knowing. *Proceedings of the Aristotelian Society*. **115**(1pt1), pp. 57–74.
- Sliwa, P. 2016. Moral Worth and Moral Knowledge. *Philosophy and Phenomenological Research*. **93**(2), pp. 393–418.
- Smith, M. 1994. *The Moral Problem*. Oxford: Wiley-Blackwell.
- Smith, M. 2004. The Structure of Orthonomy. *Royal Institute of Philosophy Supplement*. **55**, pp. 165–193.
- Stratton-Lake, P. 2000. *Kant, Duty and Moral Worth*. Routledge.

- Taylor, T. 2013. Well-being and Prudential Value. *Philosophy & Public Policy Quarterly*. **31**(2), pp. 10-17.
- Tiberius, V. 2013. In Defense of Reflection. *Philosophical Issues*. **23**(1), pp. 223–243.
- Tiberius, V. 2015. Prudential Value. In: Hirose, I. and Olson, J. eds. *The Oxford Handbook of Value Theory*. Oxford: Oxford University Press, pp. 158-174.
- Tucker, M. 2019. From an Axiological Standpoint. *Ratio*. **32**(2), pp. 131-138.
- Van Roojen, M. 2015. *Metaethics: A Contemporary Introduction*. New York: Routledge.
- Väyrynen, P. 2016. Thick Ethical Concepts. [Online]. In: Zalta, E. N. ed. *The Stanford Encyclopedia of Philosophy*. [Accessed 22 August 2019]. Available from: <https://plato.stanford.edu/archives/sum2019/entries/thick-ethical-concepts/>.
- Whiting, D. 2015. The Glass Is Half Empty: A New Argument for Pessimism About Aesthetic Testimony. *British Journal of Aesthetics*. **55**(1), pp. 91–107.
- Whiting, D. 2016. Against Second-Order Reasons. *Noûs*. **51**(2), pp. 398-420.
- Wiggins, D. 1976. Deliberation and Practical Reason. *Proceedings of the Aristotelian Society*. **76**(1), pp. 29–52.
- Wiland, E. 2017. Moral Testimony: Going on the Offensive. In: Shafer-Landau, R. ed. *Oxford Studies in Metaethics: volume 12*. Oxford: Oxford University Press, pp. 51-75.
- Williams, B. 1981. Internal and External Reasons. In: Williams, B. ed. *Moral Luck: Philosophical Papers 1973-1980*. Cambridge: Cambridge University Press, pp. 101-113.
- Wolf, S. 1982. Moral Saints. *Journal of Philosophy*. **79**(8), pp. 419-439.
- Wolff, R. P. 1970. *In Defense of Anarchism*. New York: Harper.