

A Thesis Submitted for the Degree of PhD at the University of Warwick

Permanent WRAP URL:

<http://wrap.warwick.ac.uk/160547>

Copyright and reuse:

This thesis is made available online and is protected by original copyright.

Please scroll down to view the document itself.

Please refer to the repository record for this item for information to help you to cite it.

Our policy information is available from the repository home page.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk

What's Wrong With Hypocrisy.

Kartik Upadhyaya

A thesis submitted for the degree of Doctor of Philosophy

Department of Philosophy

University of Warwick

July 2020

Contents

Abstract	5
Introduction	6
Chapter 1. Setting the Scene	18
1. Alternative Views	19
2. Hypocrisy: Worse than nothing?	29
Chapter 2. Blaming As We Should	39
1. Idealising blame	40
2. Warrant and Appropriateness.....	48
3. Engagement.....	56
Chapter 3. Non-Hypocrisy: Morally Best	68
1. Laying our wrongs ‘on the table’	71
2. Instrumentally best.....	78
3. Non-instrumentally best.....	84
Chapter 4. Non-Hypocrisy: Morally Required	93
1. Low-cost beneficence.....	94
2. Acknowledgment: cost or benefit?.....	99
3. Further worries.....	108
4. Whether versus how to blame.....	112
Chapter 5. Standing and Procedural Value	118
1. The hypocrisy case.....	120
2. Entrapment and Estrangement	125
3. Ledger and Commitment	132
4. Detachment	140
5. Acknowledgment	142
6. Equality	144
Chapter 6. Standing, on a Scale	149
1. Standing: Permission, Authority, Entitlement?.....	152
2. Standing as Appropriateness.....	165
3. Standing: Scalar by Contribution and Nature	171
4. Threshold.....	181
Chapter 7. Standing, Impairment, and Repair	185
1. Why not hypocrisy?.....	187
2. Impairment	190
3. The Reparative Project.....	197
Conclusion	214
References	214

Acknowledgments

I would hardly be ‘walking the talk’ here if I failed fully to acknowledge those who have helped form and develop these ideas. I am grateful for feedback from audiences at MANCEPT’s Brave New World Conference, 2020; MANCEPT’s Workshops in Political Theory, 2019; Liverpool Hope’s Pedagogies of Punishment Conference, 2019; Oxford’s Graduate Conference in Political Theory, 2019; Warwick’s Centre for Ethics, Law and Public Affairs, 2018; Philosophy Work in Progress Seminar, 2018; and Graduate Conference in Legal and Political Theory, 2018.

I thank Matthew Clayton, Kai Spiekermann, Laura Valentini and Mike Otsuka for support at earlier stages. For discussion and feedback on several aspects of this work over a number of years, I especially thank Matthew Clayton, Brigid Evans, Clementina Gentile-Fusillo, Simon Gansinger, Michele Giavazzi, Hwa Kim, Kasper Lippert-Rasmussen, Ida Lübben, Chris Mills, Adam Neal, Tom Parr, Fabienne Peter, Johannes Roessler, Adam Slavny, John Tillson, and Peter Wilson. For additional insightful comments, I thank Sameer Bajaj, Tom Dougherty, Johann Frick, Joe Horton, Chris Noonan, and Stephen White. The most emphatic thanks go to Kim Brownlee and Victor Tadros, for consistently attentive, perspicacious and ingenious direction on all aspects of this project, and indeed all matters related. But, I am equally grateful for their friendship, and both (yes, both!) of their humanity.

My biggest debt is to family and friends for keeping me going. I thank my dad, Prakash Upadhyaya, for constant support and concern. I thank my partner and best friend, Ciara Calvin, for sparking my interest in this topic, but mainly for her kindness, love and belief in me throughout, echoed by her entire family. Ultimately, I could never have imagined myself doing doctoral work without a lifetime of love, confidence and inspiration from my mum, Joya Chatterji, and my stepdad, Anil Seal; not to mention their immense support during the course of the thesis. Special thanks, for proofreading help, to Humaira Chowdhury.

Declaration

This thesis is my own work, and has not been submitted for a degree at another university. It contains no work published elsewhere. Chapters 1, 3, 4 and 6 draw on articles of mine that are under review.

Abstract

Hypocrisy seems to be a distinctive moral wrong. This thesis offers an account of that wrong. The distinctive wrong of hypocrisy is not a rational failing, or a deception of others. It is a problem in how we critique, and blame, others, when we ourselves are guilty of similar faults. Not only does it seem wrong to blame others hypocritically; it is also widely remarked that hypocrites ‘lack standing’ to blame. I defend both judgments. When we engage others in response to wrongdoing, there is both instrumental and non-instrumental value in ensuring that those who face similar moral predicaments reason about these predicaments, and the appropriate responses to them, *together*. Hypocrisy is wrong because it hampers our ability to realise that value. And hypocrites lose standing to blame for a similar reason. Standing depends on the value of particular people engaging in an accountability procedure together, a value which the hypocrite forestalls, by failing to acknowledge wrongdoing in the course of blaming. This account of a duty to blame non-hypocritically helps unpack the relationship between hypocrisy and a range of other defects in a person’s standing. It also responds to a central sceptical challenge to the wrongness of hypocritical blaming. The challenge is that accurate hypocritical blaming can’t be wrong because it is morally good, and indeed is morally better than not blaming others at all. My response: hypocritical blaming is wrong, not because it is morally worse than not blaming at all, but because it is second-best, i.e. worse than non-hypocritical blaming, which enables mutual deliberation about wrongdoing.

Introduction

This work originates in an argument with my partner during our first year of undergraduate study. She was advancing a view that I, trying to present myself as a nifty philosopher, was inclined to dismiss as ‘common sense’ confusion. The argument took place in the early hours, while we were searching for a drink in our communal kitchens. It went something like this:

Her: We shouldn’t steal this beer from the fridge.

Me: Why not?

Her: When the others stole from our compartment last week, we complained for ages!

Me: You’re right – we shouldn’t steal the beer. But that’s because stealing the beer would be *wrong*, not because we *complained* about others doing it.

Her: But surely the fact that we complained makes a difference?

Me: I don’t see why. We were right to complain, and we’d be right not to steal. The two things are unrelated.

Her: But it’s *hypocritical* to complain about things that we ourselves do!

Me: Maybe, but that isn’t an *additional* moral wrong. It makes no difference to whether we should steal, or whether we should complain about stealing.

Her: But surely it does...

I do not remember whether we stole the beer. But I would never convince my interlocutor about the mistake I took her to be making. Nor would I convince myself, as I now think the opposite: my scepticism was misguided, and a version of what my partner said was right. This thesis explains my change of heart.

i. Answering the sceptics

The view I was sceptical about then, and will argue for now, is that hypocrisy is morally wrong. That view is enduring and widespread; we appeal to it all the time in our social and political interactions. Whether the culprits be figures of authority in our lives, or high-profile public

officials, or people we know and love, we all get frustrated while witnessing behaviour that we deem hypocritical. What's more, this frustration is distinctly moral. It isn't just that we find hypocrisy a bit silly, irritating, or imprudent. Our outrage suggests, rather, that hypocrites behave in a way that warrants moral complaint.

Not only is it natural to find hypocrisy wrong; many of us also have pre-reflective judgments about certain aspects of the character of this wrong, two of which I will seek to vindicate in what follows. The first judgment is that there seems to be something *non-instrumentally* wrong about hypocrisy – something wrong about it, for instance, even when it causally contributes to good outcomes. We may *also* think that hypocrisy generally leads to bad consequences. But our unease about a particular instance of hypocrisy is not likely to disappear as soon as we discover that it has led to good consequences overall.

Another judgment I aim to defend is that the wrong of hypocrisy is *distinctive*: on the face of it, it seems wrong for a very particular set of reasons. Again, this tends to play out in how we respond to hypocrisy. When we complain about hypocritical behaviour, what we seem to be objecting to is precisely the *hypocrisy* involved in that behaviour, rather than something else vaguely associated with it. This suggests that hypocrisy might be wrong *as such*; it makes its own type of contribution to the permissibility of our conduct.

The aim is to give a sound argument for these pre-reflective judgments, rather than to assume that they tell us something true. Indeed, I will claim that a range of familiar ideas about hypocrisy turn out to be mistaken on reflection – for instance, that hypocrites have no right to criticise others; or that they should keep quiet rather than criticising; or that their criticisms should be ignored.

But, as a starting point, I think we should take the core idea – that hypocrisy is distinctively morally wrong – seriously as evidence of something morally important. This moral judgment is difficult to abandon. One indication of this is that the judgment animates even those who are inclined to reject it. As Kasper Lipper-Rasmussen notes:

While I know of people who deny that there is anything wrong *as such* about [hypocrisy], I do not *think* I know of anyone who does not act *as if* they believed there is something wrong about [it] – at least when they, or people they care about, are subjected to [it] (2018, p. 102).

This is true to my experience, too. I also know a few people who doubt that hypocrisy is distinctively morally wrong; yet, these doubts do not tend to manifest in their behaviour. Even these sceptics show moral distaste about others' hypocrisy, for example, as well as guilt when their own hypocrisy is pointed out; and a moral motivation to avoid hypocrisy where possible. This gives a sense of the depth of the conviction that it is morally wrong.

Whilst deeply intuitive, though, the sceptics are right to point out that the wrongness of hypocrisy is also puzzling. To illustrate the difficulties in explaining why hypocrisy might be morally wrong, consider again the view my partner described in the discussion above. One way to state that view is that there's something morally wrong, in itself, about a person doing the following combination of things: *Complaining* about a certain kind of behaviour, on the one hand, and *behaving* in a way that is morally similar, on the other. But it is not clear how this can be true, as it is not clear how either part of the combination affects the morality of the other.

We can think about this problem in two directions. First, does *complaining* about something affect our reasons whether or not to *do* it? It seems not. Consider, for example, the moral considerations we'd expect to bear on whether my partner and I should steal the beer from the fridge. Plausibly, the central reasons have to do with the entitlements and interests of the person we'd be stealing from, as well as what we'd stand to gain by stealing from them. Now suppose we add this consideration to the list: I *complained* about others stealing. What difference does this make in and of itself? It obviously would not make the difference, in and of itself, between permissible and impermissible theft: I would hardly be licenced to steal had I not complained, for example. Alternatively, would my not complaining about wrongful theft make the theft morally *better* (or, conversely, would complaining make wrongful theft morally *worse*)? That also looks unlikely. Remember that the others wrongly stole from *my* compartment in the fridge. If the others stole wrongly from me, I was surely justified in complaining about their doing so. It is hard to believe that doing something justified in the

past, in and of itself, makes doing something wrong in the future morally worse. Similarly, it is hard to see why *not* doing something justified in the past makes doing something wrong in the future morally better.

To be sure, the fact that I complained might affect the wrongness of my stealing alongside other factors. For example, the fact that I complained about a similar act may suggest that I understand that this particular act is wrong, and this suggests that I ought to have known better than to steal myself. But this is only a contingent feature of my complaint: I might complain and forget why I did so, for example. The problem is that in and of itself, the complaint does not seem to significantly affect the moral gravity of the theft.

Now think about the problem in the other direction. Can the fact that I *behave* in a certain way affect the morality of *complaining* about that behaviour in others? Again, it is not clear why that would be true. Return to the example of theft, and consider the moral considerations that would typically bear on whether we ought to complain about theft. The natural response is, 'Same as above'. The main moral reasons to complain about stealing depend, surely, on the reasons that make theft *wrong*. Complaining about this wrong voices those reasons – it voices the entitlements and interests of victims that have been infringed, and the moral implications of that infringement for the parties concerned. That is why I was justified in complaining about others stealing things from me in the first place. Now, in what way would my stealing affect that justification I have for complaining? The wrongness of *my* stealing, after all, does not affect the wrongness of *others* stealing; two wrongs don't make a right! The fact that I stole, in other words, does not alter the *accuracy* of my complaint. We can be wronged by someone even if we similarly wrong others, and we are justified, it seems, in complaining about being wronged.

To put it in another way: does the fact that I stole make it morally *worse* to complain, compared to when I did not steal, even if I am justified in complaining either way? Perhaps this is sometimes true. Perhaps when I stole, I provoked my housemates to steal, and so I shouldn't be so surprised that I ended up a victim. So perhaps my own stealing makes me partly responsible for the circumstances in which I was wronged, and in that way means my

claims over my property are less stringent than those who did not contribute to this situation; and since I am not as seriously wronged as non-contributors, my own justification for complaint may be weaker than theirs. Again, however, this would only be a contingent effect of my stealing. It is a different thing to show that stealing makes a difference *in and of itself*.

Even if we can show this, a further step is needed to show that any such difference makes complaining morally wrong. And that is difficult to do. Suppose I am choosing between two options. I can accurately complain and hold my housemate accountable for her having wronged me. Or, I can keep shtum, and not accurately complain, in which case the culprit is not held accountable for what she did. Plausibly, holding the culprit accountable is preferable to silence. Yet, this seems to raise doubts about the idea that holding the culprit accountable could be wrong. How, we might wonder, could something be wrong if it is morally better to do that thing than not to do it?

This last problem in particular drives much of the philosophical scepticism in the literature about the wrongness of hypocrisy. Were accurate hypocritical criticism wrong, the worry goes, it would be morally worse than no criticism at all. However, this is unpalatable: accurate hypocritical complaint is often clearer morally better than nothing. I call this the *Worse-than-nothing Problem*.¹ I explore it in more detail in the next chapter, and it is a central theme in this project.

ii. *A middle ground*

The task of this thesis is to show that although there is something right in these sceptical arguments, they are ultimately mistaken. There is a plausible way to understand the wrongness of hypocrisy that is not undermined by the problems raised above. Indeed, the argument I give permits most of the claims just made. I allow that complaining about something does not in

¹For scepticism about the wrong of hypocrisy in general, see e.g. Lang (forthcoming); Dover (2019); Bell (2012); Runciman (2008); Swift (2004); Statman (1997); and Turner (1990). Scepticism in the first three of these articles is driven by the Worse-than-nothing Problem last problem in particular.

and of itself make doing that thing wrong or worse. I also allow, vice-versa, that doing something wrong does not in and of itself make complaining about it wrong or worse. Nonetheless, I reject the conclusion that hypocrisy is therefore morally permissible. The duty of non-hypocrisy should be understood differently. Our own wrongful conduct does not give us a duty not to complain about others at all. Rather, it can give us a duty to complain in a particular way.

Hypocrisy, in other words, is not about our reasons *whether* to complain about others' wrongful behaviour; it is about our reasons *how* to do so – about the nature and ingredients of complaint. To complain non-hypocritically is thus to complain in the right way; likewise, wrongful hypocrisy is about complaining in the wrong way.² This analysis extends beyond the particular example of complaint, to our communication of *normative critique* in general, where this encompasses criticism, advice, demands, suggestions, warnings, and so on. In general, we ought to address these critiques to people non-hypocritically rather than hypocritically. Even if it is better to critique others hypocritically than not to critique them at all, it is best to critique others non-hypocritically, and it is normally wrong to fail to critique in that best way.

And here, in a nutshell, is why: hypocrisy disables *mutual* critique. The best and right form of critique provides a platform for similarly normatively situated agents to reason together about how to deal with normative problems they face in common. I call this view *Mutuality*. What hypocrisy hampers, wrongly, is our ability to reason together in that ideal way.

Note that whilst this view identifies mistakes in sceptical arguments, it also offers a middle ground between scepticism and the 'common sense' outrage outlined above. Although sceptics are mistaken in denying that hypocrisy is distinctively morally wrong, they are right that we should at least temper our negative reactions towards it. Outrage about hypocrisy – especially where the hypocrite is a celebrated figure – such as Donald Trump, when he

² This is not to imply that hypocrisy is always a matter of complaint. As I point out below and in Chapter 1, there are other forms of hypocrisy, but hypocritical complaint about, and more generally critique of, others is the centrepiece of this thesis.

criticises the World Health Organization (WHO) for failing to respond adequately to the threat of COVID-19,³ or Boris Johnson, when he criticises members of Union of European Football Associations (UEFA) for failing to stamp out racism in football⁴ – tends to be very loud indeed. So loud that hypocrisy can attract a good deal more attention than other serious wrongs taking place in the world. A degree of outrage is understandable, but we can lower the amplitude for sure. Albeit wrong, hypocrisy may not be as morally bad as it initially appears, because it is often better than nothing. In sum, what follows is a somewhat *deflationary* account of our practices surrounding hypocrisy and critique.

iii. Overview

It may seem striking that this introduction gives no indication of exactly what I think hypocrisy *is*, despite sketching my view about why this type of behaviour is wrong. But nor, sadly, can I claim that a conceptual definition is forthcoming. Hypocrisy seems to have several varieties, and my sense is that people will disagree on a precise set of necessary and sufficient conditions for it. Instead, the approach I take is to investigate the moral significance of different forms of hypocrisy, and of hypocrisy in critique especially. I will call all these things forms of *hypocrisy*. Ultimately we can call them what we want; but they are all in the general ballpark, and their moral significance needs examining.

Chapter 1: ‘Setting the Scene’ has two parts, the first of which explores two alternative accounts of wrongful hypocrisy, and moves the focus to hypocritical critique. The two alternatives emphasise the hypocrisy of deceiving others about our moral characters, and the hypocrisy of not practicing what we preach. I argue that relative to these alternatives, the paradigm of hypocritical critique is better-placed to meet the desiderata I mentioned earlier: to explain how hypocrisy is both non-instrumentally and distinctively wrong. I also indicate how

³ Murray (2020).

⁴ Merrick (2019)

exploring the hypocrisy of normative critique might capture what is appealing in these alternatives.

However, as suggested by the housemate example, an account of wrongful hypocritical critique has problems of its own. A central issue is what I called the Worse-than-Nothing Problem: we need to understand how hypocritical critique can wrong even when it is better than nothing. The second part of Chapter 1 explores this problem in more detail, and outlines how I propose to resolve it in later chapters.

Though what I say applies to our broader practices of normative critique, this thesis mainly discusses a specific form of it. The mode of critique predominantly discussed by philosophers in connection with hypocrisy is *blaming*. In contrast to advising, warning and merely criticising others, blaming others, as the examples of Trump-versus-WHO and Boris-versus-UEFA suggests, tends to address moral failings: serious moral failings – typically wrongdoing – which likewise require serious – even heated and exclamatory – responses. Given the seriousness of that background, and the importance of responding appropriately to wrongdoing, hypocritical engagement in this practice seems particularly objectionable and inappropriate.

One way to see this is to note how philosophers often use a specific term when discussing hypocrisy in this context: hypocrites, it's alleged, *lack standing to blame*. To see the attraction of this idea, consider a cleaner example of hypocritical blaming, to which I will be referring to throughout:

Supremacists: Amy (the 'blamee') culpably commits an act of racist abuse, witnessed by Blake (the 'blamer'). At a similar time, Blake culpably commits an act of casteist abuse. Blake blames Amy to her face for her racist abuse. Other things are equal.

Now consider this response that Amy might make to Blake: 'You're absolutely right that I'm blameworthy for racist abuse. However, you are, too, for casteist abuse. Since your casteism

is, morally speaking, a similar wrong, you lack standing to blame me.⁵ Therefore, it is wrong for you to do so, even if others may blame me permissibly.’ Beyond explaining why hypocrisy is wrong, a further aim in this thesis is to clarify this particular explanation for that conclusion: to consider what it might mean, morally speaking, to ‘lack standing’ to blame, why hypocrites lack this, and what role their lack of standing plays in making their blame wrong.

In Chapter 2: ‘Blaming as we Should’, I explain in detail what I mean by the phrase, ‘Blake blames Amy’, by giving an account of what a healthy blaming practice looks like, which also supports an account of what makes blaming apt. I argue that the core values of blaming are both communicative and procedural: We want to instil moral knowledge in the blameworthy about their moral mistakes and how to correct them; but we want to achieve that outcome through what I call an *accountability process*, where one agent interpersonally engages another on the subject of their wrongdoing, and what to do about it. The appropriateness of blaming is a matter of how it responds to these core values.

In Chapters 3 and 4, I present my argument for why hypocritical blaming is wrong. Because the main purpose of blaming is to urge *wrongdoers* to interpersonally reflect on their wrongdoing, hampering the ability of *similar* wrongdoers to interpersonally, and mutually, reflect on their similar wrongs, is all the more morally problematic.

Chapter 3: ‘Non-Hypocrisy: Morally Best’ explains how hypocrisy disables mutual deliberation between similar wrongdoers, and what we lose by disabling that process. Reasoning together about similar wrongdoing requires people to acknowledge their wrongs, yet the hypocrite fails to acknowledge wrongdoing. That failure is a value loss. Reasoning together helps similar wrongdoers advance their need to respond to their wrongs as they ought to, by combining their collective experience of the wrongdoing in question, and their shared

⁵ I say more in Chapters 3 and 5 about the importance of similarity in wrongdoing, but probably nothing wholly satisfactory by way of setting-out exactly what makes one wrong similar to another. This is a tricky substantive issue; I will tend to assume that certain wrongs are relevantly similar where they require responses by wrongdoers that are similar in stringency and nature. For a different view, and a helpful starting point, see Lippert-Rasmussen (2013, pp. 303-306).

need to figure out how to correct for it. Non-hypocritical blaming also advances our interests in jointly shaping one another's normative futures, and in the bonds of co-agency that collaboration about similar wrongdoing creates.

Chapter 4: 'Non-Hypocrisy: Morally Required' argues that this set of values is *wrongly* forestalled by the hypocrite, on the ground that mutual reflection on wrongdoing is not particularly costly for a blamer to enable. The cost of mutual deliberation is the cost of acknowledging our similar mistakes. Doing this is not too burdensome; in fact, it is often beneficial for the blamer overall. Thus, if we blame others, we owe them non-hypocritical blame, because we owe others important benefits at low cost to ourselves.

The argument in these chapters only warrants the conclusion that hypocrisy is morally wrong. By itself, this entails nothing about a further explanandum I gave above: namely, why hypocrites *lack standing* to blame. To understand the relationship between hypocrisy, its wrongness, and standing to blame, we need a substantive account of the latter. That's the business of the latter part of this thesis, in Chapters 5-7.

Providing a clear account of standing is not easy. Whilst the idea that hypocrites have defective standing is compelling, the moral meaning of this is elusive: there is no obvious grasp or consensus on what sort of moral property standing amounts to in the context of blaming; what grounds that property, or the lack of it; and what the moral implications of having or lacking it are. As I will show, people mean a wide range of things by defective standing, many of which cannot be plausible on reflection. What's more, hypocrisy is not the only instance of blaming a person with defective standing: I have defective standing to blame someone when it is 'none of my business' to blame that person, for example, or when I entrap the person I blame into wrongdoing. So, to fully explain how hypocrisy and standing relate, we need to think about hypocrisy as one among a range of possible facts relevant to standing.

In Chapter 5: 'Standing and Procedural Value', I introduce my account. Defective standing, I argue, is explained by a deficiency in the value of blaming. But that deficiency has a particular source. It is determined by what I call *connective facts* – facts about how the blamer relates to the blame, the wrong, and the fact that the wrong was committed. And the value

that these facts affect is also very specific: it is about the value of particular people engaging in an accountability process with one another. That is why hypocrites lack standing to blame – hypocritical blaming is deficient in *procedural value*, as it disables a process of mutual deliberation about wrongdoing. But a further attraction of the Procedural Value account is that it can explain a much wider range of defective cases. In general, where standing to blame is defective, it is intuitive that the reason for that defect is a procedural fault in the blame.

A crucial feature of this view is the following: The more certain facts about us enhance the procedural value of blaming, the *more* standing we have; and the more these facts diminish procedural value, the *less* standing we have. What's striking about this is that it conceives of standing to blame as *scalar*: standing-relevant facts, like hypocrisy, whether it is 'our business' to blame someone, and so on, can often make us have more or less standing, rather than giving us standing where we otherwise lacked it, or taking standing away altogether. That is at odds with ordinary discourse on blaming – ordinarily, we talk about standing in *binary* terms. We say things like: 'Blake has *no* standing given his hypocrisy'; and 'he's in *no* position whatsoever to say that to Amy', or he has '*no* right at all to complain'. However, as I show in Chapter 6: 'Standing, on a Scale', this orthodoxy is misguided, and has led to downstream confusions in our debates. Standing considerations don't in themselves give us rights, duties or permissions; their core function is to make blame more or less morally appropriate, and, as such, to guide us on whether we ought to persuade certain people to blame rather than others, and at what cost people should blame as appropriately as possible. That is because standing to blame is a matter of degree.

This overall picture – Mutuality and Procedural Value, combined with a Scalar view of standing – illuminates our assessment of more complicated instances of blame, where a mixture of different standing-relevant facts is in play. One such instance is often used as a paradigm for a lack of standing: namely, cases of *mutual wronging*, where agents blame one another for having similarly wronged one another. Chapter 7: 'Standing, Impairment and Repair' considers why mutual wronging cases attract special attention. I raise a puzzle about this case: since the hypocritical blamer is a victim of the other person's wrongdoing, we'd

expect their standing to blame to be enhanced; and yet, it seems no better than it is in non-mutual cases, like Supremacists, where the blamer is not a victim. To address this puzzle, I argue that the reason why mutual wrongdoers lose standing is that they pass up a procedural opportunity – an opportunity to ensure that future relationships with between wrongdoer and victim can proceed unimpaired. This view improves on alternative accounts of standing in mutual wronging cases, which appeal to the idea that wronging a person impairs a pre-existing relationship with them. That backward-looking approach helps explain the relevance to standing of relationships in the ordinary sense, such as being someone’s partner, friend, or colleague: the impairment of these pre-existing relationships can diminish the appropriateness of holding someone accountable. However, the same is not true of agents who wrong each other without that intimate background. The appropriateness of a wrongdoer creating the conditions for a future unimpaired relationship with his victim is not sensitive to past wronging, and accountability processes are a central way of appreciating this fact.

To conclude the thesis, I clarify how its arguments meet the various desiderata identified above, and hint at how these arguments might be further applied and developed.

Chapter 1. Setting the Scene.

This chapter explains my focus on hypocrisy in normative critique, and prefaces the structure of the argument to come.

Why take critique to be an appropriate target for a project on wrongful hypocrisy?¹ There are, after all, ways of behaving in a manner that people consider hypocritical, and morally wrong, that do not involve critiquing others. One popular view, for instance, is that wrongful hypocrisy consists in not practicing what we preach: we should ‘walk the talk’, some people think. Another prevalent view is that wrongful hypocrisy consists in a certain kind of deception; a ‘wearing of moral makeup’, or an attempt to present ourselves as morally better than we actually are. We can do these things – fail to walk the talk and deceive people about our character – without criticising other people. So why focus on critique, and why focus on it exclusively?

In Sections 1, I outline problems with both these alternative views. I don’t reject either of them. I only point out that they are not obviously true, and that they face difficulties in providing a satisfying account of what makes hypocrisy wrong. However, I also think that there is *something* true in both ideas; and I indicate how the account that I will defend – namely that wrongful hypocrisy consists in critiquing others when we warrant similar critique – might meet limitations in these alternative accounts whilst also explaining some of their appeal.

¹ Recall that by *critique*, I mean a range of different kinds of normative engagement with others, such as advice, accusations, warnings, demands, and blame. Whilst the focus will shortly move to the last of these, I’ll use this word for now to refer to the whole range – please forgive its repetition in what follows.

That is not to suggest that relative to these alternatives, the wrongness of hypocritical critique is easier to establish. Quite the contrary; the wrongness of hypocritical critique has puzzled philosophers, mainly in the light of the following observation. Hypocritical critique can be accurate: it can communicate normative truths. When it does so, it seems positively valuable – it’s better than *silence*, at least! A tempting thought is that hypocrisy is therefore not wrong, as were it wrong, matters would be opposite: it would be better not to critique people at all than to critique them hypocritically. I explore this problem, the *Worse-than-nothing-Problem*, in the second part of this chapter. In Section 2, I explain it in more detail, and outline the kind of argument that I will embark upon to respond to it.

1. Alternative Views

Some people observe hypocrisy in certain kinds of behaviour that may, but need not, involve critiquing others. Here I consider two broad accounts along these lines, and explain the advantages of a critique-centred view.

i. Walking the Talk

A familiar way to characterise a hypocrite is to say that she ‘says one thing, yet does another’. It is also commonplace to take moral issue with that sort of behaviour. This suggests the following account:

Walk the Talk: Hypocrisy paradigmatically involves the failure to practice what we preach, and that failure is its central wrong-making feature.

To illustrate this view, consider the case we started with: of stealing people’s things from the fridge, whilst also complaining about that sort of behaviour. The problem of hypocrisy, on this view, is that the behaviour doesn’t live up to the norms expressed in the complaint – if I complain, I should ‘walk the talk’ and not steal myself. Note, though, that the *complaint* about others is not actually a central part of what makes hypocrisy wrong. Complaint, and generally criticising others’ behaviour, is one way of preaching about norms. But it is not the only way. I might preach about theft loudly on the street, for example, without addressing

anyone in particular, or claiming that anyone in particular is guilty of it. According to Walk the Talk, if I then steal myself, this is no less a case of wrongful hypocrisy than in the case where my preaching involves critique of others. Whilst I think there is some insight in Walk the Talk, it also faces problems; and a focus on hypocritical critique helps us understand both the insight and the problem. I will set out a general problem, before raising concerns about specific arguments.

a. No reasons to be rational

No-one who believes Walk the Talk, of course, believes that we cannot morally assess the behaviour and the preaching in isolation of one another. It is true that the stealing is wrong anyway; but the thought is that when we combine stealing with preaching about it, the hypocrisy of this makes stealing *more* wrong. The same thing applies to preaching: preaching might be a morally unattractive thing to do, but it is made all the more morally unattractive (indeed, wrong) when a preacher herself does not do as she recommends to others in her preaching. But when we think about the type of failing that not practicing what we preach amounts to, it is unclear both whether this is a source of moral wrongdoing, and if it is, whether it is the kind of wrongdoing that hypocrisy seems to be.

One possible type of failing is a kind of *akrasia*, or acting against one's better judgment.² If I believe in the substance of what I preach, namely that stealing is wrong in certain circumstances, and then steal in those circumstances, I act against my better judgment.³ Another possible failing is what we might call *substantive* irrationality. Suppose I do not really believe in the substance of what I preach: I don't really think that stealing is morally wrong, and that is why I violate the norm myself. That may suggest that I don't fully appreciate the type of thing a moral norm is: a norm is meant to play a role in our deliberations about how

² For a view that distinguishes weak will and hypocrisy, see Szabados & Soifer (1999).

³ For more extensive discussion and critique of this particular version of Walk the Talk, see Lang (forthcoming), and Dover (2019).

we should conduct ourselves, and it is meant to play that role for all agents in similar circumstances.⁴ If I fully appreciated these features of norms, we might think, I either would not invoke this norm in my preaching without believing in it; or if I did, I would not violate it myself.

Now, *akrasia* and substantive irrationality are significant failings in an agent. But they do not in themselves seem to be *wrongmaking*. They might make a person more vicious, or less virtuous; a virtuous person is committed to moral norms, whereas both rational failings – not having a strong enough will to comply with norms, and not appreciating the sort of thing a norm is – are evidence of a lack of moral commitment.⁵ However, not everything that makes a person less virtuous is morally wrong. The reason why I am a person of dubious morals might owe to the fact that I am lazy, for example, but I don't act wrongly simply by being lazy. And it is particularly difficult to explain what is morally wrong about these rational failings as such. A plausible view about norms of rationality is there is no general reason to act rationally for the sake of it; no general reason to do what we believe we should do, for example, or to do what we intend, or to take the necessary means for achieving what we intend.⁶ Whether we have any such reasons depends entirely on the accuracy of our attitudes: on whether what we believe we should do, or what we intend to do, is actually what we ought to do. We can make the same point about the accuracy of what we preach. If, for example, I preach that stealing people's things is morally *permissible* or *required*, I have no reason to do what would be strong-

⁴ Following Michael Smith (1995), we could call these things the Objectivity and Practicality of moral judgments (1995, ch.1).

⁵ Compare Crisp and Cowton's view (1994) that hypocrisy is the vice of not taking morality seriously. We might extend this idea to non-moral normativity; part of the agential failure remains even when the norms we fail to take seriously are non-moral. But see Lang (forthcoming) for the view that there is no vice where the norms at issue are mistaken.

⁶ For arguments in this vein, see e.g. Raz (2005a) and Kolodny (2005), and for further discussion, see Raz (2005b), Kolodny (2008a; 2008b; 2005); Broome (2005); and Holton (2004).

willed, or what would show appreciation of the type of thing a norm is. That is because the preaching is morally inaccurate.

We can summarise the problem this way. Norms of rationality are *content-independent*; they apply regardless of the accuracy of our attitudes. In contrast, what we have moral reason to do is *content-dependent*. Moral reasons depend on facts, and where our attitudes or preachings are mistaken about these facts, we have no moral reason to comply with them at all. So, we should doubt that we have any general reason to walk the talk, let alone a moral reason which would make failing to walk the talk wrong.⁷

b. Who's wronged?

There may be a version of Walk the Talk which avoids that general concern. We might think, for example, that it is wrong to fail to practice what we preach *on the condition* that what we preach is accurate. On this view, when I accurately preach about stealing, it is more wrong to steal, than it would have been had I not accurately preached about it. Let us briefly consider some positive considerations we might advance in favour of an appropriately revised version of Walk the Talk.

One thought is that those who fail to comply with what they preach lack *integrity*. We commit a kind of self-betrayal, or a disrespect for our own agency, when we violate norms we endorse or invoke. Some might think that it can be morally wrong to act without integrity.⁸ But, even if this is true, it doesn't quite square with the sort of moral wrong hypocrisy seems to be. If hypocrites wrongly betray themselves by not living up to their commitments, the

⁷ Some do argue that there are general moral reasons to be rational (e.g. Bratman 2009a; 2009b). But even if this were true, it wouldn't show that it is morally wrong to be irrational, of course – perhaps being rational is generally morally supererogatory.

⁸ See Nili (2018).

primary victim of that betrayal is ourselves⁹. However, a duty of non-hypocrisy seems primarily owed to others. That feature of the wrong is well-evidenced by observations made in the introduction: when we criticise hypocrisy in others, we generally complain about it, rather than pity it – our immediate reaction to Trump criticising the WHO for a poor response to COVID-19, for instance, is not, ‘Poor you!’ It is more like, ‘Don’t interact with others like that!’ The same is also plausible when we reflect on our own hypocrisy. For me, at least, discomfort about my own hypocrisy does not arise from a sense of self-pity; it arises from the fact that others are subjected to it.

Now compare the idea that wrongful hypocrisy has to do with how we critique others when we make similar mistakes ourselves. This type of activity is naturally other-regarding: it involves one agent aiming to engage another about how she ought to behave. So there is no particular structural challenge involved in explaining how hypocritical critique might engage the interests of other people, and how critiquing people in certain ways rather than others might be a better or worse way to treat those people.

Of course, it may be possible to show that failing to practice what we preach does wrong others. For instance, we may think that people who don’t practice what they preach are generally unreliable: we have an interest in being able to plan our futures – this often requires reliable expectations about what others are going to do, and part of what forms these expectations is what others say. Hypocrites may frustrate these interests where their preaching is not a reliable indication of their future behaviour. Another reason why hypocritical preaching might wrong others is because it cheapens the value of preaching. Gerald Lang (forthcoming), for example, sketches the view that hypocrisy is problematic in a manner similar to other vices in normative speech, such as grandstanding and bullshitting. As with these other vices, Lang argues, people who do not practice what they preach are not committed to what

⁹ I think integrity is a conditional value: it is valueless for a person with evil principles to manifest these in her behaviour (compare Nili, 2018, pp. 429-432). So this argument only works on the revised version of Walk the Talk.

they're saying, with the result that over time, what people say about moral matters is taken less and less seriously. This is a bad thing; and it is bad for everyone, given that we all have an interest in moral talk being taken seriously.¹⁰

Whilst I think these arguments are importantly true, they are not entirely satisfying. One limitation is that they imply that the moral disvalue of hypocrisy is a contingent matter. Consider the argument about reliability, for example: my failure to walk the talk hardly has implications for others' plans where they don't *know* about my talk – those who don't know that I preach do not have faulty evidence about what I might do in future.¹¹ The same applies to the argument about cheapening moral talk. It seems that moral talk is only devalued when hypocrisy is exposed and known about. But again, the fact that others don't know that we violate the norms we preach about does not make preaching any less hypocritical on the walk the talk view. So the cheapening idea does not show that hypocrisy is in itself bad.

This reiterates a conviction I expressed earlier. Hypocrisy seems, at least in part, non-instrumentally wrong – its wrongness doesn't come hand-in-hand merely with whether its consequences are good or bad. But the foregoing raises another issue: the people harmed by the relevant bad consequence of the cheapening view are underspecified. If moral talk in general is cheapened over time, that is bad for everyone in general. Yet, it seems that hypocrisy often has specific victims. Though people in general might be harmed by hypocrisy in the long run, we tend to distinguish primary complainants from the general set of people harmed.

Again, compare, on this score, an account which holds that hypocrisy is a form of wrongful critique. On this view it is clear who in particular stands to be wronged: namely, the

¹⁰ Also compare Isserow & Klein (2017).

¹¹ True, people *should* be able to rely on my not stealing anyway, and should reasonably assume that I won't steal in making plans with me. But this holds quite generally, and independently of what we preach about: whatever people say, we should ideally be able rely on each other not to act wrongly. Furthermore, not all wrongdoing affects the plans that others are entitled to make, such as when I burn a precious book that no-one knows I own just for fun. This frustrates no-one's rightful plans, but, if I preach against book-burning, this is no less of a rational failing.

addressee of the critique. And it is also easier to imagine how the addressee might be non-instrumentally wronged. Addressing people normatively has an interpersonal structure: it is a way in which a person engages the capacities of another, to try to get them to reflect on a certain normative matter. The way we treat each other in our interpersonal interactions looks like the sort of thing that can have a non-instrumental significance.

Another concern about these arguments for Walk the Talk is that they do not tell us anything morally distinctive about hypocrisy. Not practicing what we preach is not the only way for people to be unreliable – compare someone who barely preaches but is just disorganised. Nor is hypocrisy likely to be the morally most troubling source of unreliability in people – consider someone who breaks promises. Nor, as Lang explicitly says (forthcoming), is the cheapening of moral talk a *special* problem with hypocrisy; what’s centrally wrong with it – namely that moral talk is not taken seriously – is exactly the same as what’s wrong with grandstanding and bullshitting.

But hypocrisy seems morally distinctive;¹² and the focus on critique promises to explain why. Critiquing others whilst guilty of similar mistakes is a very specific type of act, even a very specific form of critical engagement. And the wrongness of this form of engagement also seems specific: it is not like critiquing someone falsely, or vindictively, or in order to humiliate them. There is something uniquely troubling about an interaction where we invite someone to reflect on, and respond appropriately, to certain reasons, and fail to recognise that similar reasons apply, with similar stringency, to ourselves. Friends of Walk the Talk will point out that this failure also suggests a rational failing, or a mistake about how reasons work: reasons should play a role in the deliberation of *all* similarly situated agents, and those who critique others fail to apply these reasons to themselves.¹³ This may often be true. However, the attraction of introducing the factor of interpersonal critique is that it explains

¹² Compare Fritz & Miller (2019), Isserow and Klein (2017), and Wallace (2010).

¹³ Compare Tadros (2009, pp. 394-404) for a view that associates hypocrisy with a non-uniform application of reasons without defending Walk the Talk.

how the importance of agents appropriately recognising their similar normative situations might create a situation where one person clearly wrongs another particular person. This, I hope, will come out more clearly in what follows.¹⁴

ii. *Moral Makeup*

A second familiar view about the wrongness of hypocrisy, which connects it to deception, has a similar shortcoming:

Moral Makeup: Hypocrisy paradigmatically involves deceiving others about our moral past or our moral character, and that deception is its central wrongmaking feature.¹⁵

This account provides a different way of explaining what the relevance of my behaviour has to my preaching. When I preach about the wrongness of theft, I signal to others the type of person I am – I portray myself as someone who either has not wrongly stolen or would not wrongly steal. Yet, my behaviour (i.e. my act of theft) suggests that this is not true of me. This suggests that my signal to others is deceptive: I'm communicating something false about my character, or trying to, at least. Notice that as with Walk the Talk, critique is unnecessary according to Moral Makeup: I can present myself as better than I am by criticising other thieves, but I can do that in other ways, such as by maintaining a general saintly demeanour, as *Tartuffe* managed to do in Molière's play. Moral Makeup holds that this kind of deception is central to what makes hypocrisy wrong. Again, I think there is something right in this view, but that it is limited, and focussing on critique helps bring out both its appeal and its limitations.

Compared to Walk the Talk, the challenge of explaining how hypocrites might wrong particular others, in a non-instrumental way, is less powerful here. Part of the wrong of deception does seem non-instrumental – as with hypocrisy, we are not normally morally

¹⁴ See especially Chapter 3.

¹⁵ Versions of this view feature in Szabados & Soifer (2004); Statman (1997), McKinnon (1991) and Kittay (1982).

comfortable with deception whenever its consequences are better than non-deception.¹⁶ And the primary victims of deception are easier to specify: namely, those we intend to deceive.

Care must be taken, though, to specify the conditions in which someone is wronged by character deception. Note that such deception might be achieved by failing to reveal facts about ourselves, and we do this all the time without feeling bad about it. Politeness is an example. Suppose that though I'd never act on it, I have a morally base desire to hit the barista at the local café because I find him annoying. It is doubtful that I wrong him by not revealing this. Even if I act particularly nicely to the barista with the intention that he thinks I'd never hit anyone, he does not have a weighty complaint against me. Indeed, we arguably have powerful interests in not having to go around revealing things about ourselves, as is evidenced by strong protections of our privacy.¹⁷ So a general duty to reveal all our moral faults looks difficult to vindicate.

Again, a version of Moral Makeup may overcome this worry. Perhaps wrongful deception of others about our characters isn't so much a matter of failing to reveal facts about ourselves, as it is a matter of communicating false or misleading information about ourselves.¹⁸ If the barista suspects that I am annoyed, and asks me what it is about him that annoys me, perhaps I wrong him by lying. This is more compelling if I have already acted in way that is clearly wrong – if the barista asks me whether I stole from the café, and I did so, it seems clearly wrong not to tell him.

However, even where deception about our moral faults *is* wrong, that wrongness seems to have little to do with hypocrisy. Imagine this interaction:

Lying:

Housemate: Did you steal my stuff from the fridge?

¹⁶ For different explanations of this judgment, see e.g. Pallikkathayil (2019), Shiffrin (2014), O'Neil (2012), Faulkner (2007).

¹⁷ Compare Nagel (1998).

¹⁸ See e.g. Shiffrin (2014, p. 12) and O'Neil (2012).

Me: No.

Suppose what I say here is false, and suppose I say it knowingly, with intent to deceive my housemate. This seems wrong, but it does not look like a wrong of hypocrisy at all. It is just outright lying.

Similarly, we can have wrongful hypocrisy without deception.¹⁹ Imagine this different interaction:

Shamelessness:

Me: That was bad of you to steal – you need to return the stuff.

Housemate: Fine, but I saw you steal too!

Me: Don't be silly.

Housemate: But you *saw* me watching you steal!

Me: No chance.

Housemate: Sigh.

Assume we both know that what my housemate is saying is true. And assume we both know that we know this – I know that my housemate witnessed me stealing, and she knows that I was aware of her witnessing it. On these assumptions, it is hard to understand this as a case of deception. Still, the way I interact with my housemate is both hypocritical, and wrong.

In response, we might distinguish, as Jennifer Lackey (2013) does, between *deceiving* someone and *being deceptive*.²⁰ Though I am not misleading my housemate here, I *am* play-acting – I'm behaving as if something is false when it is true, and it may be apt to describe this as deceptive behaviour. Still, I don't think this quite explains the role that hypocrisy is playing in such interactions. Here is a slight variant of the case:

Diversion:

Me: That was bad of you to steal.

Housemate: Same for you. I saw you do it the other day!

¹⁹ This is close to what Wallace (2010) describes as shameless hypocrisy (2010, pp. 313-315).

²⁰ Lackey (2013, *passim*)

Me: Right, but you really must apologise and return those peoples' things.

Housemate: Sure, but so do you. And you *saw* me watching you steal; we laughed about it together!

Me: Stop changing the subject!

If we again assume that what my housemate says is true, I am not play-acting here. Rather than play-act, I confess to what I do. I'm not being deceptive about the issue; I am simply diverting the conversation away from it. But, again, it is very plausible that I am behaving hypocritically, and wrongly. What this suggests is that wrongful hypocrisy and wrongful deception are not morally explained solely by reference to each other. Like the faults of being unreliable, or irrational, or cheapening moral talk, the fault of wearing moral makeup can be achieved both hypocritically and non-hypocritically. What we should seek to understand (or what I in fact seek to understand) is what's special about the hypocritical cases.

Once again, the factor of critique promises to help us understand this. Notice for instance the difference between Lying, on the one hand, and Shamelessness and Diversion on the other. What is lacking in the former case, where there is no hypocrisy, and present in the latter cases, where there is wrongful hypocrisy? The difference seems to be that in the latter, I engage in critique of my housemate for what she did, whilst failing to subject myself to similar critique in the course of engaging my housemate. This factor does additional work, and suggests a way to articulate what sets hypocritical deception apart from non-hypocritical deception: it is particularly bad, *when* criticising others, to be deceptive about our own criticisability, rather than to acknowledge that uncomfortable fact. The right thing to do is to be open about our faults when critiquing others, so that we can mutually engage with the other person. Often, deception prevents that mutual engagement from taking place, though it is not the only thing that does this. I will later clarify why I hold this view.²¹

2. Hypocrisy: Worse than nothing?

²¹ See especially Chapter 4.

So much for the relative promise of hypocritical critique as a site of fertile moral discovery. Obviously, this does nothing to establish that I am right in thinking that wrongful hypocrisy consists in criticising others when we are guilty of similar criticism. In fact, one problem with this view has been a prominent source of scepticism: namely, that hypocritical critique often seems morally preferable to no critique. To finish setting the scene, let me explain the central challenge, and how I aim to resolve it.

i. Is silence better?

Suppose, for example, that Walk the Talk and Moral Makeup are both right: it is wrong to deceive others about our character by preaching certain things, and this also wrongly cheapens moral talk. This may seem to imply the following. If the preacher is not going to walk the talk, or is being deceptive about his character, he should refrain from moral talk altogether – rather than doing the hypocritical thing, he ought to keep quiet. However, it is often hard to accept that implication. The obvious instances are where the normative talk is accurate – where it accurately points out good and bad behaviour, and thus has a chance of inspiring good behaviour in others. Hypocritical moral talk, when it is accurate, is surely sometimes better than no moral talk at all.

The same problem applies, with particular force, when it comes to the mode of critique I'll be discussing hereafter: namely, accurate *blaming*. Accurately blaming a person invites that person to consider a serious moral subject – typically, their responsibility for wrongdoing. And it seems particularly good and important for a person to be encouraged to focus on that subject. Since hypocritical blaming can help a person focus on her responsibility, it seems like a good thing overall, rather than a valueless thing.

The next chapter explains at length why I think this is true. But to see the intuitive value in accurate hypocritical blaming, recall the example I gave of Blake blaming Amy:

Supremacists: Amy (the 'blamee') culpably commits an act of racist abuse, witnessed by Blake (the 'blamer'). At a similar time, Blake culpably commits an act of casteist abuse. He blames Amy to her face for her racist abuse.

Now imagine that having witnessed Amy's racism, Blake is considering two options:

- a. Keep quiet.
- b. Blame Amy hypocritically.

Which does Blake have stronger moral reason to do? Plausibly, the answer is often going to be the second thing. Despite his committing a similar moral fault, some form of Blake blaming Amy can urge her to focus on her responsibility for wrongful racism, and to respond appropriately. This may be a burdensome process of reflection for Amy to undertake, but it still seems morally worthwhile – it can be good for us, our victims, and others, that we engage in accurate critical reflection about or moral wrongs.

Keeping quiet, on the other hand, lacks these good prospects. Thus, it seems that Blake has stronger moral reason to blame Amy hypocritically and accurately than he has to do nothing. But, if it is better than nothing, how can his hypocrisy be wrong?

The worry that accurate hypocritical blaming can be morally better than silence has so far garnered three responses. But none of these are palatable. One response is that the worry is ultimately unsound, and that it is normally morally *worse* for a hypocrite to accurately blame rather than keep quiet.²² However, it is not clear why we should think this. Whatever the disvalue of hypocrisy, the moral reasons in favour of accurately blaming a person can be weighty, at least when blaming has good prospects. These reasons are not likely to be outweighed by reasons not to blame a person at all, even if hypocrisy is among them. Accurate blaming is also a morally *appropriate* response, at least in one central way: Amy is blameworthy for wrongful abuse, and in virtue of that fact, it is apt that she be invited to reflect on why what she did was wrong, and what the right way to respond to race issues in the future is. The fact that Blake himself committed casteist abuse does not necessarily prevent him from conveying this in a sensitive and proportionate way. We can be guilty of wrong things without

²² Cristina Roadevin (2018) suggests, for example, that keeping quiet is normally more justifiable than hypocritical blaming (2018, pp. 138-140).

losing the ability to understand and accurately communicate the significance of such things. This, clearly, can be better than nothing.

A second response is that hypocritical blaming is a case of *justified* wrongdoing: though there is something wrong with it, its accuracy, appropriateness and value can make it permissible overall.²³ This does not seem right either. We don't normally take the hypocrite to be 'off the hook' once we realise that his blame has some positive impact; rather, we still think he is blaming in the wrong way. Intuitively, knowing that Blake accurately conveys the facts does not 'disappear' the moral complaint against him.

That leaves a third response: when and because accurate hypocritical blaming is better than not blaming, it is not morally wrong at all.²⁴ This type of reasoning is a prominent feature in sceptical discussions of the wrong of hypocrisy. For example, in arguing against a duty of non-hypocrisy, Daniela Dover remarks:

Avoiding [criticism] would not in itself constitute a moral improvement in these cases: silence is not the answer to the [critic's] moral problems (2019, p. 405).

Another sceptic, Macalester Bell, argues:

The educational or motivational value of blame is not undermined by the blamer's hypocrisy; we can learn from the morally corrupt just as we can learn from the morally pure (2012, p. 275).

As noted at the outset, the contention that hypocrisy is permissible is also not intuitive. The judgment that we ought to blame others non-hypocritically is widely and deeply shared, and is liable to motivate even some sceptics in their everyday social interactions.

So, we may seem to be at something of an impasse. Either we say that hypocritical blaming is worse than silence, or that it is a justified wrong, or that it is not wrong at all. How do we get out of the Worse-than-Nothing Problem?

ii. The force of acknowledgment

²³ See e.g. Fritz & Miller (2018); Lippert-Rasmussen (2018; 2013); Herstein (2020); and Roadevin (2018).

²⁴ See e.g. Dover (2019), Bell (2012), and Lang (forthcoming).

A clue to the answer lies in the following observation. Suppose that in the course of calling Amy out for her racism, Blake openly and fully acknowledges his own casteist abuse. He sincerely says, ‘I know I committed that similar act; I too must work on how to respond to what I did, and I’m ready to explore this issue.’ Here, his blaming Amy is neither wrong nor hypocritical. From anecdotal experience, even those who doubt this accept that Blake’s blaming is *less* wrong, and *less* hypocritical, if he acknowledges his own wrongdoing. This suggests that one type of hypocrisy in particular might be wrongmaking: the hypocrisy (or, if you like, the additional or aggravated hypocrisy) of *failing to acknowledge similar wrongdoing whilst blaming others*. It also gives us a different way to think about the options available to Blake, and the reasons he has for choosing one over another. It turns out there are three things Blake can do:

- a. Keep quiet.
- b. Blame Amy hypocritically – fail to acknowledge similar wrongdoing.
- c. Blame Amy non-hypocritically – acknowledge similar wrongdoing.

Though the second, hypocritical option, seems morally better than the first, the third option – to blame Amy non-hypocritically – seems morally best. In other words, it is plausible that although Blake has stronger moral reason to do b rather than a, he also has stronger moral reason to do c rather than b.

This more fine-grained description of what Blake can do illuminates a way to resolve the Worse-than-Nothing Problem. In order to think that hypocrisy is wrong, we don’t need to think that it is worse than nothing; we may instead think that it is worse than a different non-hypocritical option. Hypocrisy, in other words, might be a *wrong of second-best*: though good, it is second-best to non-hypocritical blaming, and therefore wrong. On this view, the reason why hypocritical critique is wrong is not that silence is better. It is wrong because non-hypocritical critique is even better.

There are many instances of acts being wrong in this way. Take a classic example from Shelly Kagan (1989):

Parrot: I pose a high risk to my life upon entering a burning building. Once I get to the only intact room, I see a parrot in a large cage, and a large baby. I can either flee with the parrot in the cage, or flee with both the baby and parrot in the cage, but the latter makes for heavy and awkward lifting. I flee with only the parrot.²⁵

My fleeing with the parrot seems clearly wrong here, and the reason why it is wrong also seems clear. The reason is not that fleeing with the parrot is worse than doing nothing – it is better to rescue something than it is to rescue nothing. The reason, rather, is that a third option is available – namely saving the parrot *and* the baby – which is even better, and which it is wrong not to pursue in favour of saving only the parrot.

My argument about Supremacists will be similar. The wrong of Blake’s hypocritical blaming need not be explained by its contrast with inaction. The relevant contrast, rather, is with non-hypocritical blaming; *that* superior alternative is what makes hypocritical blaming wrong. And this also helps explain why it is not a justified wrong. When fleeing with the parrot, I can hardly say: ‘But look, I could have done nothing at all, and rescuing the parrot is better than nothing.’ A powerful response is: ‘But, having decided to do anything at all, you clearly ought to have done something even better.’ Since I ought to have done something better, I wasn’t justified in what I did. The same, I will argue, is true of hypocrisy: where I ought to blame non-hypocritically rather than hypocritically, I am not justified in doing the latter.

One complication for this view is that we are not always under a duty to act at all. Can we still hold that Blake acts wrongly in blaming hypocritically if it is permissible for him not to blame Amy at all? Sometimes this is true of blaming. One reason Blake might be permitted to avoid blaming is that he has already blamed people so much that he’s run out of steam. Another is that Amy is likely to lash out and harm him in response to his blame, or to distance herself from him – sometimes such risks will make silence permissible. Where Blake is morally permitted not to blame Amy at all, the allegedly wrong option (hypocritical blaming) is meant

²⁵ Adapted from Kagan (1989, p. 16), though there are similar variants in Parfit (1986). For detailed discussion of second-best wrongs, see e.g. Horton (2017), PEA Soup (2017), Pummer (2016),

to be morally superior to a permissible option (not blaming).²⁶ And it might seem especially odd, even incoherent, to think that we can have stronger moral reason to do something *wrong* than to do something *permissible*. We might think it is built-in to the idea of something being a wrong act that it cannot have stronger moral reasons in its favour than a permissible act.²⁷

However, Parrot again suggests that this is not so. In Parrot it is equally compelling that entering the burning building carries risks (to life and limb), which release me from a duty to enter the building at all. Nonetheless, *once* I decide to enter, I decide to absorb these risks. The thing to consider at that point is the benefit that the remaining alternatives can bring to others, and the costs of providing these. As it happens rescuing the baby is highly significant for her, and the extra cost of doing so compared to saving only the parrot is not significant for me. Thus, it remains not only coherent, but credible, that the second-best option is wrong, even though doing nothing is permissible.

iii. Inequality versus silence

That, at any rate, is the structure of my response to the Worse-than-Nothing Problem: I will accept the sceptical premise that hypocrisy is better than silence, yet deny the sceptical conclusion that it is therefore not morally wrong. However, it is only a structure. We are still owed a substantive explanation of *why* the options rank in such a way that hypocrisy turns out both second-best, and wrong. Why is (i) hypocritical blaming morally better than silence; why is (ii) non-hypocritical blaming morally better than hypocritical blaming, and why does the presence of this third option mean that (iii) hypocritical blaming is wrong? The next three chapters expand on these three claims.

²⁶ This relies on it being possible that it is permissible not to do something (in this case, not to blame) even when we have stronger moral reason to do it than not to do it. For cogent defence, see Harman (2016).

²⁷ See Horton (2017).

Before proceeding, it is worth briefly illustrating how *merely* appealing to the superiority of non-hypocritical blaming won't do. Consider for example Jay Wallace's egalitarian argument against hypocritical blaming (2010, pp. 328-30). According to Wallace, hypocritical blaming is wrong because it denies the equality of the blamee. The reason: we all have 'basic interests in protection from blame' (p. 329), which are equal in moral weight, and this equal weight is denied by the hypocritical blamer. When Blake, for example, blames Amy for racism, and fails to acknowledge his casteism, he subjects Amy's racism to blame without subjecting his similar wrong to a similar response. This, for Wallace, is to treat Blake's own interests in protection from blame as *weightier* than Amy's. That denial of Amy's basic equality, he concludes, is what makes his blaming morally wrong.

Wallace's argument may be able to explain (ii) why non-hypocritical blaming is morally better than hypocritical blaming: obviously, it is morally better for Blake not to deny Amy's equality than to deny it, so it is better that he acknowledges his wrongdoing, and thereby recognises Amy's equality, rather than blaming Amy hypocritically, and thereby denying her equality.

However, the egalitarian view struggles to explain why (i) non-hypocritical blaming is morally better than not blaming. It seems quite seriously bad to deny someone's equality, after all, and if hypocritical blaming *really* does this, we ought to question whether it is worse than not blaming a person at all. For just as non-hypocritical blaming does not violate equal blaming, *not* blaming also does not violate equal blaming. Thus, consider: is it morally worse accurately to blame someone in a way that denies her basic equality, than it is to not accurately blame, and thus not to deny her basic equality? Well, that seems a tricky moral question. But if it is a tricky moral question, it would not be so intuitive that accurate hypocritical blaming is sometimes *clearly* better than silence.

The point is not to give a decisive objection to Wallace's and other egalitarian accounts of the duty of non-hypocrisy.²⁸ We might state this predominant approach in a way that is compatible with what I say, and I will later discuss some contrasts and overlaps with my own view. The point, rather, is to indicate that an argument which appeals to the value of acknowledgment must be of right kind; such an argument must also explain how all three of our desiderata (i-iii) are jointly satisfiable.

Conclusion

This chapter has hinted at the type of hypocrisy I will argue against. The relevant type of hypocrisy has to do with how we respond to the fact of our own similar wrongdoing in our normative interactions with others about their wrongdoing. That kind of wrongful hypocrisy might also interact with the importance of recognising certain features of how moral reasons work, as emphasised by the Walk the Talk view, and the importance of not deceiving others about our faults, as emphasised by the Moral Makeup view. However, as will become clearer later, these alternative views do not capture the whole picture about why hypocrisy is distinctively and non-instrumentally morally wrong.

I have also sketched the sort of argument against hypocrisy that I find especially attractive: namely, one that both explains the intuitions that give rise to the Worse-than-Nothing Problem, and offers a solution to that problem.

In what follows, I flesh out that argument. I begin, in the next chapter, with an account of blaming, and what, exactly, might be valuable about it. This helps explain why hypocritical

²⁸ Similar arguments are made by Fritz and Miller (2018), Lippert-Rasmussen (2013; 2018: ch. 4); Wallace (2010); Duff (2010, pp. 127,138-9); and Tadros (2009, pp. 394-404). Although each view cashes out equality differently, for all of them, we can wonder whether accurate blaming is something that we can have stronger moral reason to do compared to the strength of our reasons not to disregard people's equality. Again, there may be a good answer, but the fact that the answer is not straightforward is a worry in itself.

blaming can offer positive value compared to not blaming, and begins to help us see how non-hypocrisy might enhance the value of blaming.

Chapter 2. Blaming As We Should.

Let's recap again the main conclusions of this thesis. First, that it is morally wrong to blame others hypocritically, and second, that we have defective standing to blame others hypocritically. Clearly, these conclusions depend on some account of the sort of activity that hypocritical blamers engage in: viz. an account of blaming. This chapter offers one.

The aim is not to give a complete conceptual account. Blaming, just like hypocrisy, is unlikely to have simple necessary and sufficient conditions. We refer to myriad things as blaming, which are sufficiently different that it is useful to disaggregate blaming into several types. Furthermore, our project is primarily normative. It is about how hypocrisy affects the morality of blaming. More important than explaining what blame is, then, is explaining why we should care about it.

The need to explain why we should morally value blame is especially pressing given one of the claims my argument is meant to accommodate: the claim that hypocritical blaming can be morally better than not blaming. When this claim is true, it is surely grounded in the value of *blaming*, rather than the value of the *hypocrisy* that the blame involves. Whilst it remains to be argued that hypocrisy can create a moral disvalue, I take it that hypocrisy can't be an additional positive moral value in itself. So, to explain the moral betterness of hypocritical blaming compared to no blaming, we need to consider why blame more generally might be a good thing. What would a healthy blaming practice achieve?

I distinguish this particular task from a recent kind of *functional analysis*: one that seeks to vindicate blame by considering the functions of blame as we currently engage in it. I also seek to vindicate blame by considering its functions, but my focus is on the functions of blame as we should ideally engage in it. A drawback of the former functional approach, I argue, is that it overlooks the question whether an alternative practice would be morally preferable to blaming. Since blaming tends to engender negative experiences in people, why not replace it with a closely related practice that is less harmful? Asking this question helps us see what blaming has distinctively to offer. I defend this methodology in more detail in Section 1.

Not everyone is a fan of blame, of course. And understandably so; the way our blaming cultures have developed will hardly incite pride.¹ However, this does not mean blame has no possible good side. In Section 2, I explore what this good side might be by giving an account of two separate criteria for assessing blame: *warrant* and *appropriateness*. Blame's warrant, I argue, is an accuracy standard: it's about being right about the facts. Its aptness, I argue, is a measure of how well its 'internal values' are realised in a particular act of blaming. To explain this view of aptness, I appeal to the idea that blaming has a basic internal structure, where one agent responds to a failing in another agent. Blame's internal values are what a well-intentioned person would minimally aim to achieve within that structure.

Appropriateness as I understand it is then a substantive standard: it is at least partly determined by what the internal values of blaming ultimately are. In Section 4, I argue that these values are both communicative and procedural. What we should aim for in a healthy blaming practice is to confer moral knowledge on those we blame; however, it also matters that the knowledge is acquired through a particular process, in which the agency of the blamee is discursively engaged. I call this view *Engagement*, and defend it by comparing it with alternative accounts of what's centrally valuable in blame.

1. Idealising blame

We use the word blame to describe things that are very different – sometimes so different that it is hard to find commonality among them. Indeed, once we look at different types of blaming response, it becomes clear that even if there is a loose common denominator, we are unlikely to discover a great deal by identifying it. So, what to do instead? One approach is *functional*: we analyse what we take to be the function of blame in paradigm examples, and argue that this function can explain why we blame as we do in many instances. Whilst I think it is fruitful to consider the functions of blame, I will adopt a more idealistic version of this approach.

¹ Compare Kelly (2018), Pickard (2011) and Franklin (2011) for discussion.

i. Disaggregation

Let's start by briefly surveying some difficulties in finding a unified view of what blaming is.

Two things we describe as 'blaming' can be quickly set aside for our purposes. I may blame the rainstorm for my bad mood, or the lightning for the burnt tree. This sort of blaming does not engage in critical assessment of the thing blamed – in blaming the rainstorm or lightning, I do not actually take these things to be criticisable, or blameworthy, for anything. That sort of judgment we reserve for agents – that is, beings with a certain level of capacity to respond to reasons. *Agent* (1) and *non-agent* blame (2) differ in that key respect. Agent-blame is the kind that concerns us, for a reason given earlier: hypocritical blaming seems to have a primary complainant, and that complainant is the blamee. The Worse-than-Nothing Problem arises from the fact that the blamee is responsible for wrongdoing, and the fact that getting an agent to focus on her responsibility for wrongdoing seems like a good thing. Only agents can be responsible for wrongdoing.

We should also not focus on *self*-blame (3). One reason is that self-blame normally involves acknowledging our own wrongdoing, and so is not normally hypocritical.² Another is that we are concerned with cases where hypocritical blaming wrongs others. *Other*-blame (4), then, is the centre-stage for hypocrisy norms.

Though we reserve the *judgment* of blameworthiness for agents, a common view is that this judgment, by itself, is not sufficient for blaming, even if it is necessary. Many argue that in deciding whether to blame a person, our interest is not merely to discover some empirical fact about them: e.g. whether that person is blameworthy. Rather this sort of fact is something we *care* about, or at least something we take ourselves to have reason to care about. We can call blame that has some normative or motivational significance for the blamer *invested* blame (5).

But although it is widely believed that blame is invested, those who believe this still disagree about the nature of that investment. One thought is that for a judgment of blameworthiness to become blame, it has to accompany a desire, for example a desire that the

² See PEA Soup (2019a) for discussion.

agent did not flout certain moral values (*conative* blame, 5.a.).³ Another way to articulate the investment of blaming is this: In judging someone blameworthy, we also take ourselves to have reasons to view our relationships with that person differently – for instance to view their trustworthiness, or spending time with them, differently.⁴ In making these further judgments about the blameworthy person, we change the attitudes that make up our relationship with that person. That change in attitudes involves a kind of distancing between blamer and blamee (*distancing* blame, 5.b.) On a third prominent ‘invested’ view, blame need not involve these fine-grained judgments about values and reasons. What basically goes on is that we get *emotional*; we get heated and upset.⁵ Such emotional reactions may be followed down the line by further attitudes, such as desires to protect certain values, or by taking ourselves to have reasons to distance ourselves from others. But that, on the emotional view, is a separate matter from the blaming itself (*emotional* blame, 5.c.). But even among those who agree that blame is emotional, there is no consensus about which particular emotions are necessary for blame. Some think it is anger or hostility at the person (*angry* blame, 5.c.i.);⁶ others think it is resentment (*resentful* blame, 5.c.ii.);⁷ still others think it is disdain or contempt (*contemptuous* blame, 5.c.iii.).⁸

However, not everyone who takes an emotional view thinks that blame is simply the feeling of an emotion; many think that the relevant emotions also have a point or purpose. But, again there is variation on which purposes are necessary. One view is that we get emotional in order to sanction the blamee – the cost of emotional attack is meant to make the person think again about behaving in that way (*sanctioning* blame, 5.d.)⁹ Another is that we get emotional in order that the blamee suffer feelings of guilt or remorse, where that suffering is

³ Sher (2006).

⁴ Scanlon (2008).

⁵ Strawson (1982)

⁶ Owens (2012); McKenna (2012a); Wallace (1994).

⁷ Wolf (2011); Wallace (2011); Bennett (2002).

⁸ Mason (2003)

⁹ Wallace (1994)

needed for them to understand the significance of their bad behaviour – why it was bad, and what they must do to correct for it (*communicative* blame, 5.e.).¹⁰

Some, on the other hand, doubt that blaming must be invested in these ways. Blame does at least sometimes seem dispassionate.¹¹ Like emotional blame, *dispassionate* blame may be private, or *third-personal* (6.1), as when we casually ‘grade’ people’s failings in their absence, either privately, or in conversation with someone else. But dispassionate blaming may also be *second-personal* (6.2) or addressed to the blameworthy person, as often seems true of judges who stone-facedly issue guilty verdicts to defendants. If dispassionate blaming is possible, that sort of blaming looks very different from the invested sort where the blamer gets morally exercised. As do third- and second-personal blaming: discussing someone’s blameworthiness in their absence is one thing; quite another is to confront that person about the matter. It is unlikely that these things have a lot in common.

Note, though, that most of these views assume rationality in blaming. Whatever invested blame involves – be it certain emotions, desires, distancing attitudes, or communicative or punitive intentions – we should allow that these things might be formed irrationally. And here is one way this might occur: we might form blaming attitudes *without* judging someone blameworthy at all. We sometimes resent and shout at people just to vent stress or anger – we know deep down that these people have done nothing blameworthy. Some think this is still a form of blaming, even if defective.¹² We might call it *scapegoating* blame (7).

This last point effectively takes us back to the drawing board. A judgment of blameworthiness looked like the most promising contender for a necessary condition of blaming, or at least of agent-blaming. Even *that* minimal consensus, though, is challenged by

¹⁰ Sliwa (forthcoming); Nelkin (2016); Fricker (2016); Macnamara (2015); McKenna (2012a); Bennett (2002).

¹¹ Hieronymi (2004); Fricker (2016); Sliwa (forthcoming).

¹² Pickard (2013)

the possibility of irrational blaming. This gives a sample of how deep disagreements have been about what blame is, and how slim the chances of successful conceptual analysis seem to be.

But there is a second concern about attempts to conceptually analyse blame. Even if there is a commonality among different types of blame, it is unclear how much we'd learn by identifying that commonality. As Miranda Fricker (2016) says:

Successful analysis delivers the highest-common-denominator set of features of X; but where X is an internally diverse practice there is a significant risk that the highest common denominator will turn out to be very low, delivering an extremely thin account (2016, p 166).

Fricker's point becomes more compelling once we are reminded that the aim in this thesis is to morally assess certain acts of blaming. Beyond being metaphysically different, the (I count at least) fifteen discrete phenomena enumerated above are also morally different. Addressing a blameworthy person, for example, has very different moral features from simply forming a judgement that this person is blameworthy. Similarly, aiming at someone's suffering will require a different moral justification from merely judging that we ought to distance ourselves from them. Furthermore, the moral appropriateness of anger at people we know to be blameworthy will surely differ from the moral appropriateness of anger in scapegoating, where we blame people without even judging them blameworthy. If we treat blame as the thing that all these discrete phenomena have in common, we will fail to appreciate these important moral differences.

ii. Now What?

Given the limited prospects and value of conceptual analysis, we would do well to seek alternative strategies. One possibility is to try to unify blame in a different way, for example by giving a *functional* analysis. On this approach, we argue that our blaming practices have a core set of functions. Though these functions may not be necessary or sufficient – they may not be aimed at or realised in every instance of blaming – they may still explain the general pattern of how we blame as we do. I likewise propose to consider the functions of blame. However, I will make the case here for adopting a more idealistic version of this approach.

Here are two recent examples of functional analysis.¹³ Fricker (2016) explores what she takes to be a paradigm case of second-personal, emotional blame, in which the blamer heatedly emphasises, to the blamee's face, the fact that the blamee has wronged her. Fricker argues that there are two key purposes in this paradigm: first, to inspire remorse in wrongdoers, and second, to inspire reform. She then argues that these two purposes can explain many cases of blame even where the relevant functions are not aimed at it. For instance, consider a third-personal case: suppose I blame some stranger who, I discover on the news, has wrongly attacked another stranger. This is not Fricker's paradigm – I know that blaming this person in this way will not inspire remorse or reform in that person. But, Fricker observes, the paradigm may nonetheless help explain it. For instance, we might think that in blaming the assailant on the news, what we're doing is saying that the *victim* of the assault is entitled to engage in the paradigm – the victim is entitled to heatedly and second-personally address her assailant, in a bid to inspire remorse and subsequent reform. Thus, the non-paradigm case is explained by functions in the paradigm case, even though the two are metaphysically different.

In a second example of functional analysis, Sliwa (forthcoming) argues that the core unifying function of blame is epistemic. The main point of the practice, she argues, is to facilitate moral knowledge of duties we incur as a result of wrongdoing: for example, duties to apologise, explain ourselves, compensate victims, and avoid similar wrongs in future. This epistemic function, Sliwa claims, has wider explanatory appeal compared to the remorse and reform functions. For instance, the epistemic function can better explain cases of blame where the victim would not be entitled to engage in Fricker's paradigm, such as where the victim 'lacks standing' to second-personally address the wrongdoer. In such instances, blame remains important as it can still facilitate moral knowledge (in the blamer, for instance, or in others parties to a blaming conversation).

¹³ There are more – compare e.g. Mason (2019) and Shoemaker & Vargas (2019), and McKenna (2012a) – but two will suffice to illustrate my point.

I won't explore the details of this dispute here. Indeed what I say ultimately resonates with both Fricker's and Sliwa's accounts: I agree that moral knowledge, reform and second-personal engagement are central to explaining why we should value blame. But notice that we can think about the functions of blaming in a more idealistic way. Rather than looking at the functions of blame that best explains the varieties of blame as we *actually* do it – we could attempt, instead, to construct the morally best version of our blaming practice. We have reason to prefer this latter approach.

One limitation of the kind of functional analysis adopted by Fricker and Sliwa is similar to the main limitation of conceptual analysis. As we saw, conceptual analyses try to understand the phenomenon of blame based on what we *actually* describe as blame. By focusing on what we tend to describe as blame, they elide the issue of which descriptions are more or less illuminating, or explanatorily important. Similarly, functional analyses aim to isolate core functions of blame based on how we *actually* tend to engage in the practice. Compared to a conceptual approach, this approach might explain a wider range of cases. But by focusing on blaming as we actually do it, it elides another important issue: namely, the functions that the *best possible version* of a blaming practice involves. The drawback of this is that it is liable to result in an incomplete view of what ultimately makes blame valuable. We should not assume that the functions which best explain our current practice – those that explain the various ways we blame as we do – are also the functions that explain why the practice is morally worthwhile.

Here is one way this can be true: the unifying functions of our actual blaming practices may not be morally worthwhile at all. Consider, for example, the related practice of punishment. Suppose we think punishment has the function of asserting power through threats of violence. This function perhaps does the best job at unification – it explains the wide range of actual instances of punishing and why we have developed particular forms of it. Yet, this particular function does not explain why the practice is morally worthwhile.

Of course, a functional methodology need not be entirely non-normative. It may be that the core functions of blaming as we currently tend to do it also happen to be morally

worthwhile, and this might be something that a functional analysis tries to explain. For example, Sliwa writes:

I suggest that we ... approach [blame] as a practice that we created in order to advance some of our interests or needs. We shed light on the practice by identifying those needs [...] it is a practice that it is overall good for us to have. The [account] thus serves to vindicate and legitimise its existence (forthcoming, p. 15).¹⁴

However, the fact that some practice is overall good for us to have, can only partially vindicate its existence. Again, consider punishment. Suppose we can unify our practices of punishment by appeal to the function of protecting people from serious wrongdoing; suppose protection explains why the practice is good overall, as although it harms and wrongs many people, the significance of those harms and wrongs is outweighed by the fact that it also prevents much more harm and wrongdoing. I doubt that showing this would fully vindicate the practice of punishment. For it is possible that we could achieve all that is good in punishment via an *alternative* practice which involves less harm. If that were so, and we would lose nothing by getting rid of punishment, we ought to prefer the alternative practice. In this sense, punishment is not fully vindicated; we have not fully explained why it is a good thing.

One reason to attempt to fully vindicate blame is that, a bit like punishment, blaming can be harmful. Blaming is the kind of thing that risks instilling negative experiences in creatures like us who get exercised about moral issues, especially wrongdoing: emotions in the blamer like those we have already discussed, such as hostility, resentment, contempt, and so on, but also in the blamee, such as shame, guilt, anxiety and so on. To fully explain what is valuable about blame, then, we must ask: what of value would be lost if we replaced blaming with something that achieves similar things without creating all these negative experiences? Replacing blame in this way is hardly feasible, of course, but asking the question helps us focus on whether blaming has anything distinctive to offer us.

¹⁴ See also Fricker's discussion of vindication (2016, pp. 165-168).

So beyond trying to unify blaming as we currently do it, a fruitful enterprise is to explain the value of blaming as we should do it, including why it is that we ought to be blaming rather than doing something else. That is the task in the remainder of this chapter.

2. Warrant and Appropriateness

i. Being (epistemically) right

One normative demand of blame is that it is warranted. But it is not clear that everyone means the same thing by this demand. Here I defend the idea that blame is warranted because (and insofar as) its content meets epistemic standards. Let's call this view *Accuracy*.

Accuracy says that blame's warrant is a matter of the epistemic correctness of the judgments it involves. If we blame someone, the judgments involved in our doing so should either be true and supported by evidence, or at a minimum, supported by our beliefs. Unwarranted blaming is inaccurate in these ways. Ultimately, I think knowledge about blameworthy conduct is what's needed for warranted blaming: if and only if the blamer knows that the blamee is blameworthy for something, can the blamer be warranted in blaming them for that thing. But thinking about warrant in terms of a broader range of epistemic correctness concerns allows us to distinguish various ways (e.g. absence of facts, absence of evidence, false beliefs) in which we can fall short of this standard.

This view helps explain our judgments about various instances and types of blame. Suppose I blame someone with no evidence that they are blameworthy, even though they are blameworthy. There is one sense in which my blame is unwarranted, and another in which it seems warranted. Accuracy explains this as follows: the content of my blaming judgments is true, so my blame is warranted relative to the facts. But it is not supported by available evidence. So it is unwarranted relative to the evidence I have.

Consider also distancing blame. Suppose I blame Amy for racism, in a way that implies that I have reason to distance myself from her. Now suppose Amy's behaviour has given me no such reasons (e.g. she wasn't racist; or she was, but has since then totally reformed and apologised to me). Or suppose his racism *has* given me distancing reasons, but those are not

the reasons I think (e.g. I think I have no reason to trust Amy, but I have every reason to), or that there is no evidence for such reasons (e.g. I have no evidence that Amy is untrustworthy in the light of what she did). According to Accuracy, these are various ways in which distancing blame is unwarranted. That diagnosis is natural.

Some might worry that Accuracy has trouble explaining what is warranted about various types of blame that lack an obvious propositional content, such as emotional or conative blame. Resentment, we might think, does not purport to be true; it is just what we feel. Similarly, our desires to protect moral values are about what we'd like to be true, rather than claims about what is true. Such emotions and desires, it might seem, can be warranted, despite not being subject to epistemic standards.

However, emotions or desires associated with blame do have propositional content in the central cases.¹⁵ For example, consider a case of scapegoating blame: suppose I resent Amy without actually thinking she made a racist remark – I'm just anxious and venting at her. This is clearly unwarranted, and a plausible explanation for why, is that the judgment that is needed to make that sort of response accurate – i.e. the judgment that Amy is blameworthy – is unsupported even by my own beliefs about the person, let alone by the evidence or facts. We can say the same about other emotions, too: I might have contempt for Amy, for example, when there is no evidence that she did anything blameworthy. Since the judgment that Amy is blameworthy is needed for contempt to be accurate, Accuracy can explain why this instance of contempt is unwarranted – it is unwarranted because the proposition that would make it intelligible is unsupported by my evidence. Similarly, I might desire that Amy not pick up pebbles. This desire is at least somewhat intelligible if I form it in response to the judgment

¹⁵ The way Gideon Rosen puts this is that there are 'thoughts implicit' in certain emotions, which can be true, false, justified, etc. (Rosen, 2015, *passim*). Also compare Macnamara (2015).

that Amy flouts certain values by picking up pebbles. But my desire is still unwarranted, Accuracy explains, as that judgment is false.¹⁶

But some might wonder whether there is more to blame's warrant than merely epistemic standards. Things seem warranted and unwarranted in different respects – saving Amy's life might be morally warranted, for instance, and saving my money might be prudentially warranted. When discussing blame, we also seem to describe treatment, rather than belief, as warranted. We often say things like, 'Amy warrants blame for her racist remark'. But, obviously, Amy is not a fact or a piece of evidence that might support a belief; she is the recipient of a certain kind of treatment. We also say things like, 'Amy's racist behaviour warrants blame', and 'blame is warranted because Amy deserves it'. Again, here we do not mean that Amy's racism, or the idea that she deserves blame, epistemically supports a belief. What we mean is that racism supports some other response or treatment.

The point of Accuracy, though, is not to deny that these phrases are coherent. The point is to isolate an independent standard – one that does not subsume under, or reduce to, others. As the above phrases indicate, warrant might describe several different things. For example, consider 'Amy warrants blame'. This may mean that Amy's being blamed is a good thing. But it might equally mean that I have reason to blame Amy; or that it is appropriate that she be blamed; or indeed that blaming her is morally justified, permissible, or required. Each idea is normatively very different, and thus ought to be distinguished, rather than assimilated into a single category. So too with the demand that blame be epistemically right: we can, and should, separate the issue of whether blame is epistemically right from whether it is more generally right.

¹⁶ A tangential issue that I'm unsure about is whether we can have warranted emotions which do not respond to any judgment with propositional content. Perhaps this is possible. We might see a person behaving violently, for example, and immediately form some blaming emotion towards them, prior to forming any cognitive attitude. If this blaming emotion is intelligible, though, it is made intelligible by a *fact* (wrongdoing) that we witness. So Accuracy may explain the intuition that the emotion is warranted.

One benefit of separating these things is that doing so allows different views to disagree about exactly which facts about a person warrant blame. Suppose I think that disgust towards someone is never permissible, or justified, or good, or apt. Then, we can naturally say that disgust is never warranted, as there are no facts the knowledge of which would make disgust the right response. If, on the other hand, disgust towards a person may be justified, apt, and so on, then contempt can sometimes be warranted, as we can have knowledge of the facts that would morally ground that response in the relevant way (e.g. by making it justified or apt).

Another benefit of Accuracy is that it allows rival views to incorporate a common internal standard. For example, suppose I think that distancing myself from someone else can be justified or made apt because of a person's bad will towards me, whereas someone else thinks that resentment can be justified where someone has wronged me. The judgments that a person has a bad will towards, or has wronged me, must then be known to be true if they are to warrant distancing or resentment.

Overall, epistemic standards matter regardless of any substantive view we take about what grounds blame. In general, we can say that *whatever* makes various types of blame good, justified, permissible, or apt, the blamer must accurately judge that those grounds obtain.

ii. *Appropriateness and internal value*

A further standard governing blame – aptness – is closely related to warrant. I will defend the view that blame is appropriate when, because, and insofar as, it instantiates, or aims to instantiate, values that are *internal* to blame.¹⁷ I will call this *Internalism*.

What exactly blame's internal values are is partly a substantive question, to be explored in the next section. The point of Internalism is that whatever they are, blame's appropriateness is about instantiating those values. Insofar as an act of blaming instantiates those values, it is

¹⁷ This echoes parts of Victor Tadros' discussion (2020a) of aptness and values in our practices of criminalisation (2020, secs. 3-5).

apt; and insofar as it fails to, it is inapt. But how exactly do we define what it is for a value to be *internal* or *external* to the practice? Well, I think we have a rough idea. Consider the value of creating people with happy lives. If an act of blame fails to create people with happy lives, that does not make it inapt; similarly, an act of blame is not made apt, or more apt, when it creates people with happy lives. Clearly, the value of creating people with happy lives is not especially connected to blaming – it is clearly external, and this, according to Internalism, explains why it is not aptness-affecting.

However, the distinction is not entirely obvious. Certain values seem closely connected to blaming, and yet have little effect on appropriateness. Consider for instance the values of helping people in general to understand wrongdoing, and deterring people in general from wrongdoing. Suppose Blake can blame Amy in one of two ways. If he does it one way, this will not help Amy, but it will help two other people, Ciara and Djamel, understand their own wrongdoing, and will deter them from wrongdoing. If he blames Amy in another way, he will help Amy recognise her wrongdoing, but will not deter or build knowledge in Ciara and Djamel. Though blaming in the first way is a better deterrent and knowledge-promoter in general, it is not intuitively more apt for Blake to blame Amy in the first way than the second. Why?

One way to get at what makes a value *internal*, in the sense I have in mind, is to think about the basic structure of blaming, and what values a blamer who puts that structure to good use would intend, at a minimum, to achieve. We can capture this by looking at how our minimal descriptions of the activity are structured: the basic form is ‘B blames A for x-ing’. This involves one agent (B) responding to a particular failure of agency (x-ing) in another particular agent (A). What would a well-motivated person making a response within that structure minimally intend? One thing might be to ensure that she herself recognises and appreciates the other person’s failure. But another thing might be that the blameworthy agent also recognises that failure: given that she failed in virtue of her agency, there was potential in her to respond to reasons in the right way, and there is potential in her now to respond to her failure in the right way. A further intention might be that these facts are mutually recognised

by *both* agents – i.e. that both appreciate, and show one another appreciation, of the importance of certain demands. I will explain these abstract thoughts in more detail below. But the point for now is that the basic internal structure of blaming only implicates two particular people – the blamer and blamee – and implicates a particular piece of conduct – the blameworthy behaviour. So, the internal structure does not require the blamer to create value for other people, or to create value in general. External considerations might require this – I might be morally required to blame someone in a way that deters lots of wrongdoing without helping that particular person – but internal values do not, as these are limited to the internal structure. This, I suggest, is why values such as general deterrence and general knowledge are not aptness-effecting.

a. Precedent

The view of aptness defended here is not meant to be particularly novel; it has intuitive appeal in other contexts. Take practices of gratitude. Plausibly there are certain values that are internal to this practice: the aspect of well-being that consists in exercising and receiving generosity, for instance, and the valuable relationships constituted by displays of generosity among participants. These values naturally account for the aptness or inaptness of particular acts within the practice. Suppose a friend gives me a keyring as a sentimental gift. Compare two responses I might make: I might throw the person an extravagant party, inviting hundreds of people. Or I might write them a private letter of appreciation. The first response seems inapt, whereas the second does not. Plausibly, the reason for this difference is that unlike the second response, the first does not fully instantiate the internal values that gratitude practices ideally achieve. The central intention involved in valuable gift-giving is not to celebrate the beneficiary publicly; it is to affirm or solidify a friendship through personal gestures.

Internalism is also similar, though slightly different, to another sense of aptness discussed in virtue theory: the idea that virtuous people act appropriately, whereas vicious people act inappropriately. As they are sensitive to values in general, virtuous blamers might be more disposed to instantiate, and aim to instantiate, the internal values of blame, whereas

vicious blamers are not, as they are less sensitive to values.¹⁸ But, the difference is that in deciding how to act, a virtuous person must consider not only the internal values of blame, but also external values and other relevant considerations. Thus, it might be virtuous to blame people in ways that are clearly inapt, as when I blame an innocent person to save thousands of lives. Whilst this may be apt overall, it is not, according to Internalism, apt *as an act of blame*.

b. Warrant: Necessary but insufficient for aptness

A final attraction of Internalism is the way it explains the relationship between aptness and warrant.

Blame's aptness is not plausibly exhausted by the accuracy of its judgments. An array of examples can show this. It seems inapt, for instance, to blame someone for committing peccadilloes at their funeral. It also seems inapt to blame someone for seriously bad behaviour in a jokey tone, and even more so where victims of wrongdoing are present. In all of these cases the blame might be accurate in its content, yet inapt. Another instance of this came out in the previous chapter: there is something inapt about hypocritical blaming, even though the accuracy of hypocritical blaming may make it better, and more apt, relative to not blaming a person at all. So, warranted blaming can be inapt. On the other hand, though, *unwarranted* blame is *inapt*. It is inapt to blame people who are not blameworthy, or whom we've no reason to judge blameworthy. Thus, while accuracy is not sufficient for appropriate blaming, it does seem necessary.

Internalism accounts for both these observations. Start with the necessity of warrant for aptness. Recall the internal structure that involves one agent responding to a certain failure in another. The value that can be achieved within this structure likely depends on accurate judgment: in order for the blamer to appreciate the significance of the blamee's failure, for

¹⁸ Indeed certain internal values may require the right attitudes in the person aiming to realise them. For example, I need to be sincere in thanking a person to achieve the full internal value of gratitude.

example, or to aim to get the blamee to appreciate this, it must actually *be* a failure, and the nature of the failure must also be known. In short, to be responsive to blame's internal values, we need to be responsive to the *facts* that ground these values, and these are clearly going to include the fact that a person is blameworthy, plus the facts that make this true. This is why only accurate blame is apt.

Internalism also explains why warrant is not sufficient for aptness. We can immediately see that the internal values of blaming are not going to be *simply* a matter of whether an agent knows that another is blameworthy for something. Knowing the fact of blameworthiness is not all that's involved in responding well to it, and in helping the blamee respond well: we must also, for example, react to the blameworthy conduct in a way that reflects the moral significance of the wrong done. That is partly why jokey responses to wrongdoing, for instance, and blaming at a funeral for peccadilloes, are inapt: these agents don't blame in a way that reflects the significance of the topics at hand.

To see the promise of this account of the relationship between aptness and warrant, compare how it builds on a prominent way to draw this distinction, explored by D'Arms and Jacobson (2002). Their account discusses fittingness: fitting responses fit the facts – they *accurately* represent their objects as worthy of those responses. Thus, laughter is fitting just when what we laugh at is funny, and envy is fitting just when what we envy is enviable. Similarly, blame is fitting just when it responds to something blameworthy. But D'Arms and Jacobson also stress that fitting responses can be inapt: laughing at something that is genuinely funny is fitting, for example, but it is inapt where laughter is offensive. So fitting blame can be inapt.

While this view has a similar structure to mine, it is less complete. Like mine, it separates the issue of whether responses are accurate (fitting) from whether they are overall apt. This explains the fact that accuracy is not sufficient for aptness. However, the fittingness thought does not fully explain why fittingness is necessary. The reason for this has to do with a more general problem with fittingness accounts: the evaluative and normative significance of fittingness is either vague or open-ended. By itself, the fact that some response 'fits the

facts' does not tell us whether or in what way we should care about its doing so.¹⁹ It is thus left open how fittingness grounds appropriateness, and how it interacts with other values. Internalism has more explanatory resources, in contrast, as it offers an account of the normative and evaluative considerations which explain aptness: a blaming response has to fit the facts because the values internal to its structure require knowledge.

To properly assess the merits of Internalism, though, we must consider the issue of what the internal values of blame ultimately are, which types of blaming do and do not instantiate those values, and how all of this matters to whether blaming is apt. I take up these tasks in the next and final section.

3. Engagement

One internal value of blame is securing corrective knowledge in the blameworthy. A further internal value depends on the way this outcome is achieved: we ought to aim to achieve knowledge by conversationally engaging the blameworthy person's normative capacities. In defence of this view, which I call *Engagement*, I explain why both these things are valuable, before arguing that they are also internal.

i. Procedural enhancement

No-one's morally perfect. It is good for each of us to recognise this about ourselves, and to be motivated by that recognition to improve ourselves, and make up for our mistakes: moral self-improvement is part of living well, and flourishing as agents, and so we have reason to care about moral self-improvement for our own sake.²⁰ Recognising our moral mistakes is also

¹⁹ This makes it difficult to assess e.g. Shoemaker's defence of anger as a fitting attitude (2018).

²⁰ See Dworkin (2011, ch. 9) for a similar idea. Like Dworkin, I don't mean to suggest that living well is only about living morally, or that living morally is lexically prior, when it comes to our well-being, to other things that also improve our well-being. The point is just that living morally is something we have a strong interest in. Also compare Tadros (2020b) and Brownlee (2019).

good for others, and it is even better that we recognise these mistakes through normative dialogue with others.

a. Doing well (by others)

One general reason why recognising our moral mistakes is good for us is that it seems generally good for us to know ourselves, and knowing our failings is part of that more general good. But, relative to self-knowledge more generally, self-knowledge about our moral failings is particularly beneficial. It seems better for me to know that I made a wrongful racist remark than to know that I have a poor dress sense, for example. We all have an interest in living well, and a key dimension of living well is living morally. One way to live morally is to live a saintly life. But, for imperfect agents like us, living morally well usually involves living morally *better* in the light of moral mistakes we make in the past. Furthermore, in order to live morally better, we need to do it for the right reasons – I do not live morally better if I correct for my wrongs because this makes me look good, for instance. To live better, I need to respond to the moral reasons for correction. Responding to these reasons involves knowing them.

Knowledge about our moral mistakes might also be intrinsically good, as it tells us something about the valence and implications of how our agency was exercised. As people's deathbed reflections often indicate, we often care a great deal about the significance of what we have done in our lives, even when we know that correction is not possible. Nonetheless, our weightiest interests in knowing our mistakes arise when there is something we can do in response to that knowledge. In those instances, knowing our mistakes can be both good in itself, and good as part of living better.

Note, though, that the strength of our interests in correcting our failings is affected by the nature of the moral mistakes we make. Thus, the stronger our moral reasons are to avoid behaving a certain way, the better it is for us to correct for behaving that way; and the stronger our moral reasons are to behave in a certain way, the better it is for us to correct for failing to behave that way. For example, if I drop something out of my window by accident, and it hits no-one on the street below, there may be little, morally speaking, to correct: I may

still be doing a good job at living well, as my mistake did not involve a failure to comply with strong reasons. By contrast, where my mistakes involve clear moral defects, or even wrongdoing, I have stronger moral reasons – indeed duties – to correct for them. So, if I drop things out of my window onto the street because I think this is funny, I have stronger interests in reflecting on what I’ve done than if I did so by accident – the latter kind of mistake reflects a serious problem in appreciating, or being motivated by, strong moral considerations. Since failing to correct such defects is a more serious failure to live well, and correction requires knowledge, the strength of my interests in correction track the gravity of wrongdoing.²¹

And in a central class of cases, our moral mistakes harm or wrong other people, as when I drop things out of the window because I thought this would be funny, and it also injures someone down below. In these cases, not only do I face a general moral pressure to correct my ways; I owe additional duties of correction *to* the people I have wronged. These are the sorts of duties I mentioned in discussion of Sliwa’s view – apology, compensation, an explanation of what happened, and avoiding similar mistakes in future. Failing to correct in these ways means both acting wrongly and wronging particular people. As not giving what we owe to victims of wrongdoing is an especially serious failure to live well, that kind of corrective failure is one which we have especially strong interests in avoiding.

In these central cases, moreover, moral self-knowledge is not merely an *aim* of the duties of correction I owe; those duties depend on it. I cannot deliver what I owe to those I wrong without knowing why it was wrong, and why that wrong creates the particular duties. Unless I understand the reasons to apologise to the person I hit, for example, I cannot fully perform the duty of apology to that person. Similarly, I arguably do not fully perform my duty of compensation to this person unless I understand why I owe it.

²¹ This is not to say that we should overall morally prefer that someone makes less morally bad mistakes. I should overall prefer that I wrongly injure someone compared to accidentally killing someone, for example. But, for my sake alone, I should prefer the latter.

These considerations reveal a further value in knowing about our moral mistakes. Not only is it good for us to know these mistakes; we owe it to others to know them. This is partly true of others in general – people in general have interests in not being wronged, which wrongdoers’ knowledge about their wrongs can help protect. But in cases where we wrong others, the interests of these particular others, i.e. our victims, in our knowing our mistakes are especially strong; strong enough to ground powerful claims on us. To be sure, victims, especially of very serious wrongdoing, sometimes care little about whether those who wrong them come to have corrective knowledge; sometimes they understandably want nothing to do with these offenders.²² But I think even victims who do not want interaction want the culprit to know that they did something wrong, and why this was wrong. Knowing this about the culprit, and knowing moreover that she has reformed in the light of her moral understanding, will alleviate a good deal of a victim’s frustration. But even if it doesn’t, victims, for their own sake, have reason to value wrongdoers’ correction, given that everyone has reason to value getting what they are morally owed.

Beyond the mere fact of getting what they are owed, victims have also interests in further implications of getting this. For example, they can have interests in being materially compensated for wrongdoing – if someone injures me it is good for me that this person covers my medical expenses. Wrongdoers’ knowledge of their mistakes also advances these material interests where this knowledge helps motivate them to compensate, and where duties of compensation are difficult to enforce.

Moreover, victims and wrongdoers alike have interests in their relationships with others, and in the quality of these relationships – interests which wrongdoers’ corrective knowledge also helps protect. Wronging others generally threatens our relationships with them. For example, wrongdoing can destroy personal and intimate associations, both

²² I suspect this is normally due to countervailing factors. Even those who are averse to interacting with the perpetrator for various reasons would probably appreciate *some* value, for them, in the perpetrator’s interacting with them in order to apologise.

instrumentally, by making people resent and distance from each other, and non-instrumentally, in the way that wrongly betraying a friend, or cheating on a partner, might diminish the value of that relationship by itself. Furthermore, wronging a person can also have implications for relationships with non-intimate people: if I wrong a stranger, it is plausible that any future relationship I might then have with this person will be impaired unless I correct for my wrong.²³ This fact is something that victims (and wrongdoers) have reason to care about for their own sake. But, again, the point is that this good can only be fully instantiated through corrective action that involves knowledge of the reasons to engage in that action. Relationships, after all, are constituted not only by the acts that people perform in response to wrongdoing, but also by the attitudes that they have towards one another.²⁴ Wrongdoers can only begin to develop the attitudes needed to repair relationships if they understand the implications of wronging the other. Relational repair will not succeed simply by saying the words, ‘sorry’, or simply by paying back for the harm caused; we need in addition to grasp our reasons to say and do these things. That grasp is corrective knowledge.

b. Collaboration and co-agency

We have explored several values of corrective knowledge about our moral failings, both for ourselves and others. That is the first part of Engagement. I will now argue that there is a particularly valuable way to achieve this knowledge in others: namely, through a process of *accountability*, or normative discussion between agents.

To bring out this further value, imagine the following example:

Self-correction. Amy wrongly makes a racist remark about Ciara because she thought this was funny. She then acquires complete knowledge of why what she did was wrong, why she

²³ I give a fuller account of this in Chapter 7.

²⁴ This is an idea of e.g. Scanlon’s (2008), again whose view about relationship impairment I will challenge in Chapter 7.

ought to apologise and compensate to Amy, and avoid similar wrongs in future. Other things are equal.

What Amy does is good; indeed, the corrective knowledge she acquires couldn't be better. Nonetheless, things could plausibly be better overall: they could be better had Amy's corrective journey involved participation in a normative discussion with another agent about her wrongdoing.

Two reasons for this are particular to this sort of case. First, whilst Amy knows her reasons to apologise, she can fully comply with those reasons only if she raises the subject of her racism in conversation with Ciara. I stressed that knowledge is necessary to comply with duties of apology. But it is not sufficient: Amy cannot purely internally apologise; she must communicate the apology to Ciara. But there is more to Amy's duty to converse than that. She cannot communicate her apology to Ciara and then swiftly *walk away*. As Linda Radzik describes the ideal form of the interaction between the parties, it involves a special "moment of recognition", marked by a "[Wrongdoer]-sees-[Victim]-seeing-[Wrongdoer] dynamic" (2011, p. 596).²⁵ That can't be achieved through a unilateral expression of apology. A dynamic between the pair that requires Amy to be conversationally open; she must be ready to listen to considerations that Ciara wishes to raise with her on the subject of his racist remark, and to be publicly responsive to such considerations.

This is a longhand way of explicating the idea that wrongdoers must be ready to be *accountable* to those they wrong. There is a loss in value when that accountability process is bypassed: the loss that consists in failing to make up for wronging others. This value can only be fully realised in a conversational format.

Here is a second reason, also particular to this sort of case, to prefer, in addition to Amy gaining corrective knowledge, that she also be open to a discursive process of accountability. Where Amy's corrective knowledge is acquired through normative discussion with another agent, it results from the exercise of her moral capacities, rather than by (say)

²⁵ See also Duff (2018b).

fluke or manipulation. This fact makes us more confident that Amy responds to her wrongdoing for the right reasons, which, as I argued, is key to our interests in living better.

Though part of the story, these arguments for engaging Amy do not generalise to all behaviour which tends to make up the subject matter of blame. For instance, imagine:

Self-reflection: Amy wrongly kills a wild goat, because she thought this would be funny. However, she does reflect seriously on what she did, and then acquires complete knowledge of why what she did was wrong, and why she should avoid similar wrongs in future. Other things are equal.

Amy's behaviour here does not create corrective duties to engage in conversation with particular others – even if she wrongs the goat, she cannot interactively apologise to the goat. And the process through which she achieves her knowledge is reasons-responsive – though internal, it arises from reflection rather than manipulation or fluke. Still, an important value of accountability intuitively remains missing from this case. Accountability's value is not exhausted by its being part of a debt to victims, and helping ensure a person responds to wrongdoing for the right reasons. It has more general value.

Here is one argument for this more controversial view. In general, when our corrective futures are jointly shaped – that is, shaped through the interpersonal exercise of our normative capacities – the outcome in which we gain full corrective knowledge is made in one way better. One reason why this outcome is better is that in it, the knowledge we gain has involved a contribution by each of us to a broader, and attractive, normative project. Call this the *Collaborative Project* of building an ideal community of agents.²⁶ In this community, agents realise their special characteristics to their fullest extent. Among these is their capacity to recognise and respond to considerations given to each other in interpersonal exchange. This particular capacity – to recognise and respond to reasons we exchange with each other – is most fully realised when agents reflect on their behaviour together with others, rather than

²⁶ Scanlon (1998) famously describes this ideal, though we needn't be contractualists to endorse it (Tadros 2017, pp. 34-35).

reflecting on it in isolation. It is this capacity of Amy's and others which is not fully exercised when she corrects only on her own: her agency, and that of others, would have been more fully exercised had others participated in the move towards correction. Thus, one kind of loss in value owes to the fact that contributions to the Collaborative Project are weaker than they could be. Collaboration in shaping a person's normative future constitutes a more valuable future.

Here is a second argument for preferring an accountability process in Self-Reflection. The process of discussing our mistakes with others realises a valuable relationship; namely a *bond of co-agency*. Addressing a moral issue with others is a way to publicly recognise and express one another's shared moral capacities, and to recognise and express the fact that as agents, we are connected in a fundamental way: we share the ability and need to recognise and respond to reasons. In that sense, we are all co-members of a community of beings who all face the same fundamental task. It is a further loss in value that no such process of recognition is achieved, and it is only achieved through an accountability process.

As Michael McKenna (2012a) says about conversational blaming, "it is part of a practice that sustains the bonds of moral community" (2012a, p. 169). We can read the two arguments just advanced as particular explications of that general thought. In combining our moral capacities to help improve one another's lives, our co-agency becomes intermeshed in a way that connects our normative futures to each other. And the process of combining our agency in itself creates a bond of mutual recognition.

These arguments are further developed in the next two chapters. They prove central to the argument against hypocrisy: hypocritical blame, I will argue, wrongly hampers this set of accountability values. However, note that the present point is not that Amy's victims, or other agents, act *wrongly* in not participating in normative dialogue with Amy. Another aspect of the Worse-than-Nothing Problem, recall, is that blaming can be supererogatory; it can be permissible not to blame a person at all. The fact that blaming can be good does not mean that it is morally required, and there is no reason to assume that because conversationally blaming someone is valuable, that people have a general duty to conversationally blame others.

The point is only the first thing: that interpersonal engagement of our agency enhances the value of corrective knowledge.

Note further that constraint: that corrective knowledge *is* ultimately acquired. It is not an implication of the agential values defended here that they should be traded-off against the value of corrective knowledge. It might be that Amy's opening her up for conversation with others makes her worse at correction compared to when she is left to correct on her own – in that case, better not to converse. My thought is only that this is less than ideal: in the ideal outcome, beyond correcting, Amy contributes to the Collaborative Project, and creates bonds of co-agency with others.

This may initially seem doubtful. Indeed, some might have the opposite view: better that people correct themselves 'off their own bat', as it were, rather than through other people's intervention. We might also worry that it is over-intrusive to have a practice of constantly involving ourselves in the lives of others, who are good at reforming on their own, purely for the sake of advancing certain agential values. However, both these worries can be met with the same response. The full value of accountability only arises between willing collaborators. We can only achieve the right kind of relational co-agency, and shape our futures jointly, if, as parties to normative discussion, we are *both* ready and willing to listen and respond to considerations that others bring to the table. So engagement, as I understand it, it is no intrusion; nor does it undermine personal motivation to correct our mistakes. Instead it is a process whose value depends on parties' commitment to scrutinising their mistakes.

ii. Pills

So much for the value of knowledge gained through an accountability process. Why think that this set of values is *internal* to blaming? I do not have a knockdown argument. However, I think these values intuitively capture what is missing in various alternative accounts of the distinctive moral contribution of blame.

One view, for example, is that blame's distinctive contribution is motivational – the point is to instil in people desires or intentions to behave differently in future. However,

imagine we could achieve this by administering an otherwise-harmless *motive pill*, without wrongdoers knowing about this. The pill implants in the wrongdoer a strong and robust desire not to repeat the behaviour. Though this achieves the same outcome, this does not fully instantiate the internal value of blaming. At least, it seems that the aptness of administering motive pills is very different from the aptness of blaming a person. What is missing?

We might think that what is missing is a sanctioning of Amy, or at least a desire that Amy be sanctioned, for her racist behaviour: the motive pill instils desires without imposing any costs, such as an unpleasant experience, on the wrongdoer. There are various possible versions of the view that sanctioning is an internal value. One is that it is particularly valuable to impose certain costs on wrongdoers because these are deserved (McKenna 2012a; Nelkin 2016). However, this doesn't seem to explain what is distinctive about blaming. Again, suppose we had a practice of administering otherwise-harmless *sanctioning pills*, without wrongdoers knowing – these pills impose certain costs on people, such as unpleasant feelings, whenever people behave badly. Suppose this gives wrongdoers what they deserve. This is not an instance of internally valuable blaming. Something is again missing.

A different view that ties the value of blaming to sanctions emphasises the emotional nature of sanctions, both for blamer and blamee – namely that it involves levelling certain emotions towards the blamee, such as anger or resentment.²⁷ Here, for example, is a reason Wallace gives for why that practice might make a distinctive contribution to our lives:

we care about [others' negative emotions towards us] not merely because it is disagreeable to experience them, but because we do not wish to inhabit a social world in which such attitudes are harboured toward us, regardless of whether they are expressed or experienced by us. It is bad for us, in its own right, to be the target of such attitudes, and this helps us to understand why it makes sense to think of them in relation to an informal but effective system of social pressures (2019, p. 92).

It's true that we care a great deal about how others regard us, and that concern surely helps motivate us to avoid and correct for bad behaviour. But, again, making people reform through an aversion to social disesteem doesn't seem *internally* valuable. Suppose we give Amy an *esteem*

²⁷ Compare Wallace, (2019; 2011; 1994) and Fricker (2016, pp. 171-76).

pill, which instils in her the knowledge that others have negative attitudes about her in the light of her racism. This may be a particularly good way of motivating Amy (and others) to change her ways. But, it is unlikely that the aptness of administering the sanction pill is explained in the same way as the aptness of blaming.

One explanation of the difference between pill-giving and blaming is the following. While we want people to be motivated to avoid bad behaviour, even being willing to impose costs of social disesteem to ensure that they do, the structure of blaming is such that it seeks to achieve a certain causal condition on corrective motivation. That is *knowledge*: as sketched above, what we ideally aim at, in blaming others, is that they respond to their failures for the right reasons. We have other practices – penalising people, for example – to disincentivise them against wrongdoing. But instilling knowledge is something that the structure of blaming in particular can additionally contribute.

Why not think that something very different from blaming, such as a gentle fleeting comment, or a grading discussion, or a general non-accusatory chat about morality, would instil knowledge equally well? In these alternative practices, negative emotions would be less intense, and the knowledge gain might be just as good. Part of the answer is that we communicate knowledge through the delivery of what we say: typically the right way to communicate the urgency and stringency of corrective demands a wrongdoer faces, as well as the gravity of wrongdoing, is to express these facts in a sombre or severe tone.²⁸ The delivery of blame may be part of what makes it unpleasant for the blamee; but, as indicated in Chapter 1, the unpleasantness of blame is often going to be outweighed by the need to convey the right message – we need to convey the duties wrongdoers have to reflect on what they did, and typically an emphatic delivery is what’s needed to get that across.

But another part of the answer is that the communication of knowledge is not the only thing that is internal to blaming. We can by now guess how I would illustrate this: we might give someone a *knowledge pill*, which instils in wrongdoers both knowledge of corrective

²⁸ Compare Pickard (2011) on ‘affective’ and ‘detached’ blame (2011, p. 218).

duties, and a good sense of their urgency and stringency. Once again, this seems valuable, but not quite in a way that captures the value of the internal structure of blaming. And we can again guess what I believe is missing: *engagement*. Blaming is not merely a transfer of content, but a discursive, interpersonal exercise of agency. No instance of pill-giving can be a substitute for this process. That is why I think that, together with knowledge, engagement captures the internal value of blaming.

Conclusion

So much for when and why I think blaming is valuable and apt. Let us now shift our focus to instances of defective blaming, starting with the defect of hypocrisy. On the face of it, notice that there is no reason to think that hypocritically blaming a person cannot achieve the internal values I have outlined. It is not impossible for people who are guilty of similar mistakes to communicate corrective knowledge to people they blame. Nor is it impossible for those guilty of similar mistakes people to shape others' corrective futures jointly with them, in a way that recognising their co-agency. Thus, it seems straightforward that hypocrites can blame in a way that is both good and apt.

Our original puzzle, then, remains to be addressed. *Even when* it is valuable, and internally valuable, hypocritical blaming also seems *inapt*, and morally wrong. The next two chapters explain how this can be true.

Chapter 3. Non-Hypocrisy: Morally Best.

The previous chapter accounted for why blaming has something distinctively good to offer in its ideal form, even if we often blame badly. I raised doubts that blame's affective components (such as the anger and resentment it involves) and its sanctioning components (such as the deterrence effect of our need for esteem by others) explain its distinctive value. Its distinctive value derives, rather, from the aim of getting the blameworthy to respond to their mistakes for the right reasons, and the means of achieving that aim is a process of engagement where we invite the blameworthy to reflect on their mistakes, to offer these reflections to us, and to consider what we have to say in response.

I also hinted at how this can explain why hypocritical blaming is better than not blaming at all – the reason is that hypocritical blaming can achieve much of the above. There is no reason to think that those who blame us while failing to acknowledge their own similar mistakes are incapable of conveying corrective knowledge to us, or incapable of engaging us in a way that invites us to reflect on our mistakes, or indeed incapable of exercising their agency with us, so that we can jointly build our corrective futures, and recognise our shared normative capacities. Recall, for example, *Supremacists*, where Blake fails to acknowledge his casteism whilst calling Amy out for her racism. Must the fact that Blake made a mistake similar to Amy's stop him from blaming Amy successfully – from engaging her in a normative assessment of her racism? That is doubtful. It is not as though by making mistakes ourselves, we thereby lack the abilities to know why those mistakes are wrong, and to reflect on such issues in conversation with others. Similarly, must the fact that Blake *fails to acknowledge* his casteism in this process stop him from getting Amy to respond to her wrongs for the right reasons? Again that is doubtful. It's not as though we cannot know and discuss the corrective demands created by certain wrongs unless we are cognisant of having committed those wrongs ourselves, or indeed unless we communicate that cognisance. To emphasise that Amy has a duty to apologise to the victim of her racist remark, for example, Blake need not add that he too must

apologise to the victim of his casteist remark. After all, the substance of such matters – *why* certain acts are wrong and the moral implications of committing them – has nothing to do with *who* is guilty of these acts. In short, hypocritical blaming seems capable of instantiating the internal values of blaming because the failure to acknowledge similar mistakes is perfectly compatible with the possession of sound moral knowledge and conversational ability. So, it seems that hypocrites can blame in a way that is both good and apt; hypocritical blaming can have a positive moral valence, and thus can be better than not blaming at all.

At this point, the Worse-than-Nothing Problem re-enters the fray. If hypocritical blaming is better than not blaming – if hypocrisy is compatible with what’s largely good and apt about blame – then how can we also think that it is *inapt*, and morally wrong? This becomes more puzzling where the option of not blaming is morally permissible. It seems peculiar, if not unfair, that morality condemns something that is not only largely good and apt, but also something that we can permissibly avoid doing, and something that is better to do than to avoid doing.

In Chapter 1, I offered the skeleton of a view that can respond to the Worse-than-Nothing Problem. On this view, accurate hypocritical blaming might be both better than not blaming and morally wrong if the following is true: the reason why hypocritical blaming is wrong is not that it is worse than silence, but rather that it is worse than the best form of blaming – in particular, worse than *non*-hypocritical blaming. According to this response, the view that hypocrisy is wrong does not require that hypocritical behaviour has a negative moral valence overall. Instead, we can show that the morality of some hypocritical act is objectionably *deficient*: its moral quality is worse than it could, or ought, to be – lower than that of a non-hypocritical counterpart. As I also noted in Chapter 1, the charge of hypocrisy in the context of blaming is most powerful when the blamer fails to acknowledge her similar mistake. If, instead, Blake shows full recognition of his guilt of casteism in the course of his conversation with Amy about her racism, he either avoids hypocrisy altogether, or avoids a significant type of hypocrisy. So, to resolve the Worse-than-Nothing Problem in the way I propose, we must explain the moral superiority of non-hypocritical blaming. We must explain, that is, why blame

which involves acknowledgment of similar wrongdoing is morally better than blame which does not. That is this chapter's task. The explanation I will give is this: compared to blame that fails to acknowledge, blame that acknowledges wrongdoing enables a morally richer type of blaming conversation – one that involves a process of *mutual deliberation*, in which similar wrongdoers can reason together about what both of them should do about their moral mistakes. It is particularly good for similar wrongdoers to engage in that mutual process. The value of this process is what we lose in blaming others hypocritically rather than non-hypocritically.

In Section 1, I clarify in more detail how acknowledgment makes possible the kind of mutual deliberation I have in mind, and what exactly this sort of conversation can achieve over and above what hypocritical blaming can achieve. Acknowledgment enables mutual deliberation, I explain, because unless we acknowledge our wrongs, they cannot be treated as a possible subject of discussion – a subject that is on the table for joint scrutiny with those we blame. Joint scrutiny requires a two-way conversation about both parties' wrongs, in contrast to hypocritical blaming, which allows at most a one- or two-way conversation about only one of the parties' wrongs.¹

In Sections 2 and 3, I give three reasons why we have stronger moral reason to enable that additional conversational possibility. One reason, explained in Section 2, is instrumental: compared to hypocritical blaming, mutual reasoning between similarly placed wrongdoers tends to help them to engage in the right corrective responses to their wrongs, by further building their knowledge and motivation to respond appropriately. But as I contend in Section 3, mutual reasoning between similarly placed wrongdoers also has non-instrumental value. One such value is *constitutive*: compared to hypocritical blaming, mutually deliberative blaming involves a fuller exercise of both the blamer and blamee's normative capacities, which, I will argue, is part of the best outcome for both parties. Another non-instrumental argument, a

¹ A one-way conversation might seem impossible. I mean something like Amy 'talking at' Blake with no response. I will clarify 'two way street' I have in mind below.

relational argument, holds that mutual reasoning between similar wrongdoers makes possible a special bond of solidarity between them.

These ideas will be familiar from the previous chapter: ultimately, the values that I take to be advanced in non-hypocritical blaming are those of corrective knowledge and engagement. But, the point is that the kind of engagement available in non-hypocritical blaming is more valuable than the kind of engagement available in hypocritical blaming. I conclude by explaining how this helps us to see why hypocritical blaming is inapt in one sense, despite being good and apt in another.

1. Laying our wrongs ‘on the table’

Why, in the first place, should we think that acknowledging similar wrongdoing when we blame others enables mutual deliberation about how to respond to our similar mistakes? The basic thought is that refusing to acknowledge wrongdoing takes it ‘off the table’ for discussion, whereas mutual deliberation requires both parties’ wrongs to be ‘on the table’ for discussion. In *Supremacists*, for example, Blake and Amy may be in a position to discuss Amy’s racism. However, because Blake does not acknowledge his casteism, the pair are not in a position to discuss *that* issue; and this latter fact prevents any deliberation from being *mutual*. The aspect of a blaming conversation ruled out by hypocrisy, then, is the ability for interlocutors to assess one another’s wrongs on a two-way street, or a multilateral platform. Before spelling out why that is an important loss, let us explore how hypocrisy creates it.

i. Wide evaluative scope

The finer details of what mutual deliberation between blamer and blamee involves will, I hope, emerge as we proceed. But to introduce the type of communicative interaction we might expect in non-hypocritical blaming, I will draw on an idea of Johann Frick’s (2016). Frick describes a model of justification – what he calls ‘interpersonal justification’ – between two agents, A and B. On this model, B addresses A as follows:

‘A, you should x in circumstances C, because your x -ing follows from a sound argument about what *we* (i.e. A *and* B) should do in C.’²

To illustrate, consider an example of Gerald Cohen’s (1992). Suppose Bea is a kidnapper, and Alf is a parent of a child Bea has kidnapped. Bea is unwilling to let the child go for free, and she is trying to persuade Alf to pay a ransom. Note, in these circumstances, that Bea is *not* in a position to justify to Alf why he should pay the ransom in the way described above. Bea cannot accurately say to Alf: ‘Alf, your paying the ransom follows from a sound argument about what we (you and I both) should do’. To the contrary, Alf’s paying is ruled out by that consideration. What Alf *and* Bea should do involves Bea returning the child to Alf for free, and if Bea does this, Alf has no reason to pay her a ransom.

To be sure, there *is* a genuine justification for Alf’s paying; the fact that Bea will not release the child unless Alf pays is a compelling reason for Alf to pay. And this is a reason that some implicated parties can cite to interpersonally justify to Alf why he should pay: a police officer could cite this reason, for example, as could Alf’s child. But, the point is that *Bea* cannot cite that reason to the same end. Unlike the police officer or child, Alf’s duty to protect his child is engaged only because Bea is herself unwilling to do what that duty demands, having placed the children under threat. So, whilst the child’s protection is clearly a reason, this reason does not follow from a sound argument about how all parties to a normative discussion between *Alf and Bea* ought to act.

As the term ‘interpersonal’ has broader connotations, let us rename the kind of justification that is unavailable to Alf a *mutually deliberative* justification. As we can see, to be in a position to give a mutually deliberative justification, what we need to do is widen the range of things we are willing to evaluate when engaging with others. That is, we must draw a

² Frick’s focus is different, however. He considers whether Scanlon’s contractualist theory can explain cases of hypocritical demands, rather than hypocritical blame, or why it is wrong.

conclusion about what others ought to do by publicly evaluating the normative situation of not just others, but also ourselves.

At first glance, it might seem odd to apply this line of thought to Blake and Amy in Supremacists. The deliberative situation between them is not remotely like a hostage negotiation. The relevant circumstances (C) are the parties' mistakes, and the relevant response (x-ing) is corrective – what Blake ought to be engaging Amy about is how to respond to her moral mistake. Unlike Alf the parent, who only has reason to pay because his addresser, Bea, acts wrongly; Amy, the blamee, must correct for her racist abuse regardless of whether her blamer, Blake, corrects for his casteist abuse. As I pointed out in Chapter 1, two wrongs don't make a right, and obviously, two failures to correct for wrongdoing do not make a right either. Thus, unlike Bea to Alf, Blake does not have to revise his own intentions in order to give Amy a sound justification for how she must ideally respond to her wrongdoing.

However, the question here is not whether mutually deliberative justification is needed to justify a certain course of action to Amy. Our question is what is needed for Blake to be in a position to offer such an argument in the first place. And in this respect, there is a similarity between Blake the blamer and Bea the kidnapper: whether these agents can give their interlocutors a mutually deliberative justification depends on their willingness to do something. For Bea, it is her willingness to release the child unpaid. For Blake, it is his willingness to acknowledge his mistake. Let me explain.

ii. *Conversational self-exclusion*

Why would a mutually deliberative justification on Blake's part require acknowledgment of wrongdoing? To see this, consider an example of how such a justification might go:

'Amy, you ought to apologise and reform in response to what you've done, because your doing so follows from a sound argument about what *we* should do in response to what *we* have done'.

As with an argument made by a kidnapper who is willing to release the child, an argument of this sort involves public appreciation of how both agents (Blake and Amy) have acted, and how that wider evaluation yields certain normative implications for both agents.

Blake, clearly, is in a position to offer such a justification if he is ready to acknowledge what he did. Where he openly recognises, and is ready to reflect on what he did, it is possible for him to treat this fact as a premise in a moral argument that he could publicly offer to Amy. It is possible, in other words, for him to openly reason from a claim about what both of them did to a claim about what both of them should do, and to communicate that reasoning to Amy. In this sense, acknowledging wrongdoing lays our wrongs on the table; it treats them as a possible conversational subject.

In contrast, Blake is *not* in a position to give Amy a mutually deliberative justification if he does not acknowledge what he did. When he is not openly acknowledging of his casteism, that wrong is no longer a possible premise in an argument he could publicly give to Amy. To see this, consider two things that an unacknowledging Blake *can* do, which might seem similar, but are not in fact the same. One thing an unacknowledging Blake can do is allow Amy to *know* the fact that Blake too ought to apologise in a certain way, by giving Amy an argument that implies something about what he himself should do.³ For example, Blake can say, ‘Amy, you ought to apologise to the victim in a particular way, as this sort of apology is apt for people guilty of supremacist abuses’. If what Blake says here is accurate, the argument he gives implies that he himself ought to apologise in this way to the victim of his casteist abuse – and there is no reason why Amy, provided she knows what Blake did, cannot draw this conclusion herself. Another thing Blake can do is *entertain* a mutually deliberative justification, and even state such a justification to himself. There is nothing to stop him from internally widening his evaluative scope to consider himself as well as Amy; so there is nothing to stop him from grasping and reciting a mutually deliberative justification ‘in his head’, so to speak. Blake can *privately* acknowledge wrongdoing, we might say, without *publicly* acknowledging it.

³ Thanks to Joe Horton for discussion.

Granted, these things are possible for Bea the kidnapper as well. Bea might pretend to be a police officer, for example, and say something like, ‘Alf, it would overall be wise to pay, though obviously you and your child are being wrongly treated in being asked to pay.’ Saying this implies that Bea acts wrongly and may imply that Bea knows she acts wrongly, which Alf also knows if he sees behind her pretence. Whilst that implication is not something she is not ready to communicate directly or be moved by, it may still be something she and Alf are both able to internally understand.

However, private acknowledgment of our duties, and even saying things that allow others to know about our duties, is not the same as giving someone else a mutually deliberative justification. What we need, in addition, is to treat those duties as a possible conversational subject: a subject that we can address *together* with those we blame. It is in this sense of ‘conversational self-exclusion’ that the failure to acknowledge forecloses genuine mutual deliberation between blamer and blamee.

iii. Two-way Streets

I will argue in the rest of this chapter that this foreclosure of mutual reasoning also means a loss in value. But before proceeding to those arguments, let me explore in a bit more detail what distinguishes mutual deliberation from other modes of deliberation.

Frick himself does not focus on this question. However, he hints at an answer:

[Mutually deliberative] justification, it might be said, is how we get one another to do things qua moral agents (Frick, 2016, p. 262).

Read appropriately, I think this passage begins to get to the heart of the matter: as the passage suggests, the distinguishing factor of mutual deliberation has something to do with how it advances and recognises the *agency* of deliberators. But, how exactly should we understand that loss?

The point cannot be the following: to get blamees to respond to certain considerations ‘qua moral agents’, we must blame them in a mutually deliberative way; so hypocritical blaming fails to treat the blamee as an agent. Whether Blake acknowledges his own wrongdoing or not,

he gives Amy a sound argument for what she must do about her wrongdoing. Provided he does this, he engages her moral capacities, and thus treats her as a moral agent.

We might instead hold the converse view, that in order to treat *ourselves* as moral agents, we must blame others in a mutually deliberative way; so hypocritical blaming fails to recognise the blamer's own moral agency. But this is not quite the point either, as it does not appropriately explain what's *wrong* with what Blake does. The wrong of hypocritical blaming, as mentioned in Chapter 1, seems 'other-regarding' – the main complainant is the blamee, not the blamer. Yet, Blake's denying his own moral agency seems bad only in a self-regarding way.

Note how both these accounts are incomplete in opposite ways. The first identifies a genuine agency-interest of Amy's, but one that is not frustrated by Blake. The other identifies an agency-interest that is not an interest of Amy's, but which can be frustrated by Blake. What is missing from both?

Here is my suggestion: each account tells us a half of the full story. The difference made by Blake acknowledging his wrongdoing is not merely that it enables him to engage *Amy's* normative capacities; nor is it merely that it enables him to engage his *own* normative capacities. The point of mutual reasoning, rather, is precisely its *mutual* quality. A genuinely mutual blaming interaction involves the parties *combining* their normative capacities multilaterally to address the full range of pressing issues: namely, both of their similar wrongs, rather than only one. Only then can the pair reason about their wrongs *together*, in the intended sense, as opposed to reasoning in isolation.

By 'combining our capacities multilaterally', I don't mean anything like merging ourselves into a collective agent. What I mean is basically open discussion about a distinctive subject matter. There is a sense in which we can combine our capacities to scrutinise a set of subjects without much discussion at all. For instance, I can raise a set of subjects with you, and each of us can go away and think about it at home. But, openly discussing an issue allows us to combine our capacities in a more *co-responsive* way. It means offering others ideas we take to be relevant to some issue, expecting them to listen to and consider those ideas, and to offer any responses to what we have to say. But it is not just 'one-way traffic': we are also ready,

ourselves, to absorb, reflect on, and respond to considerations that the other person wishes to raise, including considerations they have raised in response to what we have said. Described in detail like this, mutual deliberation about wrongdoing might seem long and complicated, but it is nothing more than what we routinely do in willing conversation.⁴

But we should remain careful about exactly what kind of joint scrutiny is ruled out by hypocrisy. Clearly *some* joint scrutiny is possible in hypocritical blaming too. Imagine that, in the course of his entire blaming interaction with Amy, Blake refuses to acknowledge his casteism, even when pressed. That in itself, does not stop the pair from jointly scrutinising a number of things. For example, Blake can offer Amy several thoughts on the subject of her racism, and on how people generally ought to correct for supremacist wrongs. He can even say, ‘Amy, people who do this type of thing should apologise’. And he can be responsive to many things Amy might say about this issue. For instance, if she asks, ‘how should people who do this type of thing apologise?’, he can say, ‘one appropriate way is intimately and face-to-face, say by inviting the person to a one-to-one meeting.’

What the pair cannot jointly scrutinise, though, is precisely the thing that Blake does not acknowledge: namely, his own supremacist abuse, and his need to correct for it. When it comes to *that* issue, Blake cannot participate in a co-responsive interaction. He *can* listen to and think about what Amy has to say about it. Suppose Amy says, ‘you made that degrading casteist joke; you, too, should apologise to your victim as well.’ Blake can fully grasp that point internally. But a discussion of it cannot be on the table, or on the two-way street that I described, unless he acknowledges his wrong. Without acknowledgment, he can neither respond to Amy’s reflections on the matter, nor can he offer his own critical reflections. He

⁴ This develops the idea in the previous chapter that a blameworthy person’s duty is to be ready to answer to a victim, instead of just walking away. Of course in this case, the blamer (Blake) is not a victim of wrongdoing – so we need to think more carefully about what other kinds of value a co-responsive interaction between Blake and Amy will add.

cannot e.g. say, ‘apologising to the person I insulted needs a different tactic’; or, ‘I will apologise to the person, but she does not want to hear from me right now...’

We can now spell out more intelligibly the ideal of mutual deliberation between wrongdoers that hypocrisy prevents. The ideal requires that parties to blame be ready to reason together about a certain range of subjects: namely, *both* of their similarly culpable acts. It is only when the full range of these subjects can be reasoned about together that a deliberative process is genuinely mutual. The reason why hypocrisy prevents that ideal is that the possible range of subjects of two-way reasoning is limited: only the wrong committed by the blamee can be openly discussed; the similar wrong committed by the hypocritical blamer cannot. And this prevents interlocutors from exercising their normative agency in the richest possible way, i.e., in a co-responsive, or multilateral, way.

Let’s take stock. I have argued that a certain kind of hypocrisy prevents similar wrongdoers from achieving the ideal of mutual deliberation: the hypocrisy of my being unwilling to publicly self-evaluate whilst evaluating you, prevents us both from having an open discussion about my wrong, and thus prevents us from having an open discussion about both our wrongs.

But so what? Why should we care morally about having fully engaged in mutual deliberation; in what sense is it a better agential exercise? As I have conceded, there is still plenty to discuss with hypocritical blamers; and regardless of whether we are blamed hypocritically, we might have ample opportunities to discuss our wrongs with others who do not blame us, or who do blame us but are not culpable for wrongs similar to ours. What is the distinctive moral advantage of mutual deliberation between wrongdoers about both of their similar wrongs?

In what follows, I give three responses to this question. Mutual deliberation adds instrumental, constitutive and relational value.

2. Instrumentally best

In Chapter 2, I argued that one internal value of blaming is corrective knowledge: to help the blamee understand what they need to do in response to what they did, and act on that understanding. In this section, I argue that mutual deliberation between similar wrongdoers advances this particular internal value of blaming – it helps promote correction in both parties to blame by improving corrective knowledge; and improving motivation, both to respond appropriately to wrongdoing, and to blame people in a proportionate way. It won't always be helpful in the ways I describe, but I will nonetheless argue that at least some of the benefits are likely, and so should be taken seriously in the decision whether to acknowledge similar wrongdoing.

i. Epistemic benefits

Here is one way in which mutual deliberation between wrongdoers about both their wrongs (as opposed to only one of their wrongs) can enhance knowledge of their corrective duties:

Transfer. As similar wrongdoers, mutual deliberators may have special performative knowledge about how someone can come to commit an offence like the one they have committed; first-hand experience of what makes such offences wrong;⁵ what it feels like to be guilty of them; and how best to set about responding to them. A process of mutual deliberation can help one party convey such knowledge to another, thereby improving the other's moral understanding of her wrongdoing and how to respond to it.

Where this benefit applies, the moral understanding of one party is limited in some respect, and this limitation is remedied through the other interlocutor's conveying their knowledge. Blake's experience of making a casteist remark, for example, might enable him to understand some key aspect of that type of wrong – say, how to respectfully approach victims to apologise to them. When he is ready publicly to assess that experience with Amy, he can

⁵ Compare Bell (2012, p. 273). Also compare Brownlee (2015), who argues that certain virtues, such as *aspiring* to be good, are especially well-known by the morally faulty (2015, pp. 5-8).

draw on it to illustrate to Amy what an appropriate apology requires in practice.⁶ Admittedly, Blake could convey the same information without acknowledging his own wrongdoing, but acknowledgment allows him openly to appeal to his own experience to make his points more credible and vivid.

Consider now a second epistemic benefit of mutual deliberation:

Joint Acquisition: A process of mutual deliberation can lead to both interlocutors improving their moral understanding of their wrongs and how to respond to them.

Where this benefit applies, both interlocutors' moral understanding is limited in some respect, and this shared limitation is remedied through the insights that the deliberative process helps both parties to discover. Suppose, for example, that Blake and Amy both lack proper understanding of how to apologise for racism; let's say that both their attempts to apologise have failed. Through reflecting on and diagnosing one another's failures together – seeing where each of them did well, and what went wrong – the pair can both end up with an improved sense of how to address victims. Again, we should note that this is not necessarily going to happen: perhaps, prospects are better when we speak to someone else who succeeds in apologising. But, remember the options we are morally comparing. We are assuming Blake is going to speak to Amy anyway, and we want to know what way of doing this is more likely to be knowledge-enhancing. Plausibly, the prospects in terms of knowledge, either through Transfer or Joint Acquisition, are generally going to be better than the prospects of these things if Blake refuses to reason about his wrongdoing together with Amy.

ii. Motivational benefits

As well as not guaranteeing epistemic benefits, mutual deliberation between wrongdoers is also not necessary for achieving them. Blake might successfully transfer the relevant knowledge, and he and Amy might jointly acquire knowledge, even if they are engaged in non-

⁶ This is only likely to happen where Blake has learned from his mistakes, or is willing to learn. But, that willingness is part of the ideal of mutual deliberation I presented earlier.

mutual deliberation. Blake might say, ‘Amy, my friend did something similar and here is how it played out...’, and the pair might learn from discussion of that more anonymous description of events. Often, though, the transfer and acquisition of knowledge will be more effective where the blamee sees that the blamer has first-hand experience of the wrong in question, and is open to reflecting on that experience. Knowing that about those who blame us will often make us more attentive to what they urge us to do and to reflect on.

This points to a further benefit that mutual deliberation can bring – it can better motivate wrongdoers to engage in corrective reflection and action in the following way:

Listening: Interlocutors in mutual deliberation can be more willing to take seriously what each other say when they show each other recognition that they are in the ‘same moral boat’.⁷ To illustrate, suppose Amy has a certain set of defensive attitudes towards those who blame her for her racism – e.g. others haven’t been ‘in my shoes’, or are singling me out, or are trying to feign moral superiority – and is thus less responsive to being called out. Blake’s willingness to openly discuss his wrongdoing helps her to overcome those psychological hurdles. He shows that he has been at least in similar shoes; that he is not singling Amy out, as he is ready to subject himself to similar scrutiny, and that he cannot be claiming moral superiority, as he fully acknowledges that his own moral record is similarly tainted. With these worries allayed, Amy can better absorb the content of the blame.

There is a further instrumental benefit of the fact that mutually deliberative blaming makes it hard to act morally superior. Not only does this help motivate the blamee to listen to the blamer; it also helps motivate the blamer to blame more appropriately:

⁷ There is some supporting psychological literature on this. As studies of the phenomenon of ‘Do-Gooder Derogation’ show, people who are otherwise morally responsive are more likely to resent and dismiss criticism when they believe their critic is playing the moral high ground. See e.g. Minson and Monin (2012), Cramwinckel et al (2013) and Monin et al (2008).

Proportionality: Mutual deliberation helps ensure that the tone of blame is sensitive to the gravity of the wrong in question, and that the harms of blaming are justified by the importance of the blamee's need to correct for that wrong.

As suggested in the previous chapter, communicative blame is sometimes too gentle in its tone to get its point across. But we often have the opposite problem of *overblaming*, where blame involves gratuitous harm and exaggerated outrage. The hostility directed at Amy may be harsher than the level necessary for her to understand and respond appropriately to the message about her racism, or indeed harsher than is appropriate given the stringency of the duties she has to correct for it. Mutual deliberation mitigates this problem.

One reason people overblame is that they come to believe that what the culprit did is either morally worse or more blameworthy than it actually is. That is much less likely in non-hypocritical blaming: when Blake openly discusses his similar guilt, he and Amy can remind him of just how bad prejudiced abuses are, and how bad a person is who is guilty of such abuses. This naturally makes it harder to exaggerate how blameworthy Amy's abuse is.

Another reason why people overblame is to *grandstand* – to use blame of others as a means of presenting ourselves as morally exemplary. And, it should be clear, that issue also disappears in non-hypocritical blaming: there is little point in Blake grandstanding in front of Amy when his own culpability is under their joint scrutiny.⁸

Now, whether blame is well-absorbed, and delivered in a proportionate tone, is a separate issue from whether it ultimately leads to reform in wrongdoers. One thing is to be motivated to listen to and reflect on what others say; another is to be motivated to *act* correctively after a blaming interaction. However, mutual deliberation can be additionally helpful in that respect, as follows:

⁸ It might also be noted that hypocrites are especially prone to grandstand in the first place, as suppressing our own errors can distort the way we think and talk about others' errors. On broader concerns about grandstanding, see Tosi & Warmke (2020; 2016).

Confidence. Interlocutors might be more motivated to comply with their corrective duties (for the right reasons) through mutually deliberative dialogue.

Suppose that Amy struggles to believe that she can change her racist ways, or that the burden of reform worries her, making her reluctant to try to change. Now suppose that, through evaluating his experience of casteism with her, Blake could explain to Amy how someone similarly situated has been, or is, able to critically assess, or revise, their prejudices, and how this might be a relieving and rewarding thing to do, rather than a burden.⁹ That sort of interaction might give Amy reason to think that reform is a worthwhile endeavour within her reach, and thus strengthen her motivation to pursue it.

iii. Contingency

As I have been wary to keep flagging, how likely mutual deliberation is to bring any of these corrective advantages, is ultimately a contingent matter. In some cases, the best corrective outcome can be perfectly well-achieved when we reason about our wrongs in isolation, or with non-similar wrongdoers, or with saints, rather than with similar wrongdoers. Indeed, sometimes, mutual reasoning between similar wrongdoers may make it harder to reform – even if parties are willing to openly self-scrutinise at the beginning of the process, they may sometimes end up rationalising and solidifying their similar flaws, finding villainous allies in one another.¹⁰

But, this does not show that there is no moral reason to give it a go: after all, some of our moral reasons are set by what we can ideally achieve, rather than what will actually happen. I also think it is armchair-psychologically plausible, and plausible in the light of readers' own experience of dialogue about wrongdoing, that our readiness to reason about our wrongs when

⁹ Compare Edwards (2018, pp. 24-25).

¹⁰ Moreover, although I have cited some psychological evidence that claims to moral superiority over others makes them reluctant to listen to us, we cannot confidently infer from this that non-hypocrisy will; and overall, I suspect the effects of mutual deliberation would be difficult to accurately test.

criticising similarly situated agents (and indeed agents who are on the verge of being similarly situated), is typically a highly effective corrective tool. And there are widespread practices to suggest that those familiar with the psychology of serious wrongdoing are all too aware of these benefits, such as the practice of inviting convicts and ex-convicts to address young offenders, in schools, prisons, and juvenile detention facilities.¹¹

Furthermore, the value of mutual deliberation does not depend only on its causal contribution to the right corrective outcomes. Even where it makes no causal difference to the likelihood of reform, it has other kinds of value to offer. So I argue now.

3. Non-instrumentally best

Beyond being instrumental to reform, the process of mutual deliberation has constitutive value, because it is part of the best corrective outcome. It also has relational value, because it creates a special type of bond between deliberators.

Let's start by noting that blame without acknowledgment seems lacking in value even where blamer and blamee are already well-poised to respond appropriately to what they did. This comes out in the following example:

Sage Supremacists: as in Supremacists, plus the following. Amy already knows why her racist abuse is wrong, and knows which corrective demands her wrong creates. Moreover, Amy knows that Blake, too, knows why his casteist abuse is wrong, and knows which corrective demands it creates. Both are motivated to act on that knowledge. Nevertheless, when blaming Amy for racism, Blake fails to acknowledge his casteism.

Here the pair reason impeccably in isolation: they do not need each other's help at all in correcting for their wrongs. Nonetheless, some improvement on the way Blake blames Amy is intuitively possible and called-for; and this seems at least partly explained by the absence of a mutually deliberative interaction. Whilst it is good that they reform well by reasoning alone,

¹¹ Dakers (2011) is a useful discussion.

it is intuitively better for them to reason together, and better still to have reformed well *by* reasoning together. Why?

i. Collaboration+

In Chapter 2, I argued that in ideal blaming, the value of corrective knowledge is enhanced when achieved through a process of engagement. Reform is better when achieved via an interpersonal exercise of agency, where agents jointly shape one another's normative futures. I motivated this idea by offering the ideal of a Collaborative Project, in which we seek to build a community of agents all of whom shape their conduct by exercising their normative capacities to their fullest extent. The thought was that the more our agency helps build one another's futures, the more moral value those futures have in virtue of that fact. This idea applies even more forcefully to mutual deliberation.

I said that what distinguishes non-hypocritical blame is the platform it offers for parties to combine their normative capacities in a multilateral, or co-responsive way; to listen to and respond to another's views about two pertinent subjects (both interlocutors' wrongs) rather than one (the blamee's wrong). This involves a fuller exercise of the parties' agency, on matters that are central to both of their lives. Thus, non-hypocritical blaming allows for a stronger contribution to the Collaborative Project than is possible in standard, unidirectional blaming.

When a subject can be multilaterally-reasoned, we know that more agents use their agency than they might otherwise do, as we know that two people rather than one are considering that subject. But that is not the key point here, since agents can still *consider* subjects by reasoning in isolation. The key point is that in multilateral reasoning about a subject, we also use a wider variety of agential *capacities*. Not only, that is, do we use the capacities needed to normatively reflect on a subject; we also use the further capacity for what I called co-responsiveness – i.e. to absorb and reply others' ideas, to consider their replies to our ideas, and so on. That is the sense in which multilateral reasoning involves a 'fuller' exercise of agency

than unilateral reasoning – agency is more fully exercised to address a normative topic when more types of agential capacity are deployed in addressing it.

Again, let's be careful: the point is not that *only* non-hypocritical blaming can achieve this type of value. Indeed, I have suggested that this value is something that blame in general can and should aim at. And the value of collaboration is possible even in hypocritical blaming. In hypocritical blaming, Blake and Amy can collaborate about *Amy's* wrongdoing, and thus can build her responses to her wrongdoing together.

However, this is second-best. It is *even* better that the pair can collaborate about both Blake's and Amy's wrongs, and thus build both their normative futures together. And that is only possible in non-hypocritical blaming. So, non-hypocritical blaming is like *Collaboration+*; it allows that two topics be multilaterally addressed, rather than one. It allows more collaboration, and so involves a stronger contribution to the Collaborative Project than hypocritical blaming.

To illustrate the appeal of the idea that outcomes can be better the more collaborative they are, compare a related thought in political philosophy. A prominent view about democratic legitimacy is that the ability for citizens to willingly participate in political decisions adds value to those decisions. One argument for this is constitutive: willing participation in elections and public debate constitutes an ideal type of political community. These types of collaboration make the outcome better even if the same decisions would have been made without citizens collaborating at all.¹²

A natural extension of this thought is the following: the *more* political issues citizens get to collaborate on, the more value the outcome has; so the more issues they collaborate on, the closer they get to the ideal of a jointly-shaped political community. My argument for collaboration about wrongdoing has a similar structure. Non-hypocritical blaming constitutes a more valuable outcome as it allows collaboration on more topics (again, on both parties'

¹² The literature on procedural legitimacy is enormous – see Peter (2017) – but the closest analogues might be Leland (2019) and Lister (2017) and Ebels-Duggan (2010).

wrongs, rather than just one of their wrongs). This comes closer to the ideal of a jointly-shaped normative community, and thus adds value even if wrongdoers would reform without any collaboration at all.

ii. Solidarity+

Some might wonder, though, whether that added value is particularly relevant to non-hypocrisy. It is not as though Blake and Amy are *uniquely* able to collaborate – either of them could collaborate with a third party, and either could ensure that the other does so. This can happen even if Blake blames hypocritically. To bring this out, imagine that Blake has another option beyond (a) silence, (b) hypocritical blaming and (c) non-hypocritical blaming: (d) he could blame Amy hypocritically before getting Ciara involved in the discussion, who is not a wrongdoer at all, but is a willing conversationalist ready to reason with Amy about a number of shared concerns. Is d not at least as good as c? If Blake chooses d, he can ensure just as much collaboration, only the subject of collaboration cannot be his casteism, at least not with him among the collaborators. So, the value of collaboration does not fully explain why non-hypocritical blaming is a particularly important type of collaboration. Here I offer a further explanation: a special kind of solidaristic bond is available between similar wrongdoers through reasoning together about their wrongs – a bond that is not available in reasoning together with other random people.

It might seem counterintuitive to suggest that it even *could* be morally preferable, in itself, that wrongdoers reason with other wrongdoers rather than reasoning with rightdoers, or indeed that wrongdoers reasoning together is better than rightdoers reasoning together.¹³ Compare, for example, the following worlds:

Sinners: Amy and Blake reason together about what they should both do in response to their racist and casteist abuses.

¹³ Thanks to Matthew Clayton for discussion of this worry.

Saints: Saint 1 and Saint 2 reason together about what they should both do if confronted by racist and casteist abuse.

Other things equal, it seems absurd morally to prefer Sinners over Saints. How could a world in which sinners are reasoning together be any better than one in which saints are reasoning together? In fact, the opposite seems true: Saints involve no wrongdoing, whereas Sinners involves wrongdoing, and surely worlds that involve wrongdoing are morally better than those that do not.

However, the value of wrongdoers reasoning together does not imply that it is better that wrongdoing exists, or that we have reason to ensure that wrongdoing exists in order to enable wrongdoers to reason about it together. The thing to contrast Sinners with is not Saints, but rather these other possibilities:

Sinners: Amy and Blake reason together about what they should both do in response to their racist and casteist abuses.

Saints Only: Amy and Blake commit racist and casteist abuse. The saints reason together about what Amy and Blake have done, but Amy and Blake do not.

Sinners with Saints: Amy and Blake commit racist and casteist abuse. Amy reasons together with Saint 1 about this, and Blake reasons together with Saint 2 about it.

My thought is that, in one way, Sinners is morally better than both Only Saints, and Sinners with Saints. To put it another way: *Given that similar wrongdoing exists*, it is particularly valuable, in itself, that similar wrongdoers reason about this subject together. When put this way, the idea is not absurd, but intuitive. Unlike Saints, racists have a morally urgent subject to consider, and they have that subject in common: namely, their responsibility for similar wrongdoing. It is thus better that they, rather than others, reason about this subject, and it is even better that they reason about it together.

Here is an argument. When similar wrongdoers, like Amy and Blake, reason about their situation, they relate to one another in a valuable way. Reasoning together about their similar situations involves public mutual recognition of the fact that they are co-agents, both morally capable, and involved in a common endeavour of responding to reasons. This creates

a solidaristic bond between the pair: each recognises that they and their interlocutor are not alone in the business of responding to certain reasons. Rather, they're in it together.¹⁴

Now, I argued in Chapter 2 that bonds of co-agency are internally valuable in blame more generally, and that such bonds are possible between all agents. Saints can recognise this with each other, and Amy and Blake can recognise this with Saints too: all of them are co-agents, and reasoning together is a way for all of them to recognise that they are all in the business of responding to moral demands. But the solidarity achieved in those cases is not quite the same. The demands that Amy and Blake confront are both *urgent* – they are duties to respond to serious wrongdoing – and *morally similar* – they are duties to respond to wrongdoing of the same type. These facts connect Amy and Blake more deeply; and the value of them mutually recognising deeper connections enhances the bond of co-agency that they can share.¹⁵

This is not to say Amy cannot recognise with a Saint, for example, that the same reasons *apply* to the Saints as they do to her. That is a general fact about reasons – they apply uniformly to all agents in the same circumstances, regardless of identity. Saints can say, ‘Amy, *were* I in your circumstances, then I would be required to apologise, reflect on, and so on’. That is one type of agential solidarity. But, it is another thing to openly recognise that we actually *confront* the same type of reasons as the people we are interacting with – that we are actually both *in* the circumstances where the relevant reasons apply. That deeper kind of solidarity –

¹⁴ I suspect that the value of this is not simply in avoiding the bad feeling that we get in thinking that we are alone in a normative endeavour, though this may also be significant. On the severe harms and injustices associated with loneliness, see Brownlee (2016; 2020).

¹⁵ One way to make this intuitive in the real world is to think about ex-offenders and current offenders. Consider a prison programme that encourages wrongdoers to reflect on their wrongdoing. One policy involves getting ex-offenders, who have now been released, to talk with the prisoners about the wrongdoing. Another involves getting random good people to do this. Other things equal, the former programme seems appropriate; the latter not so much! A plausible explanation is that a certain kind of relational value consists in the former interaction, and not so much in the latter.

Solidarity+, if you like – is a further value missing in hypocritical blame relative to non-hypocritical blame.

This helps explain what I think is both importantly true and incomplete in the Walk the Talk view described at the beginning, which conceives of wrongful hypocrisy as a kind of rational failing. We can now see how a certain kind of rational failing – a failure to publicly appreciate the fact of confronting similar moral demands – might matter to our moral assessment of hypocrisy. It matters for the types of bonds we can create with others that we recognise normative facts we confront in common, and how in virtue of similar normative situations, we confront similar demands. While this matters, notice that it is not the same as avoiding irrationality. To avoid irrationality, we can simply recognise the relevant reasons *privately*: we are not being irrational so long as we believe that certain reasons apply to us as much as they apply to others. But hypocrites often know this all too well, as Sage Supremacists suggests. The problem with hypocrites is that they prevent the solidaristic bonds with others that this knowledge makes possible, as these bonds depend on our similar situations being not only known and understood, but also openly and mutually recognised and appreciated.

The value of these bonds also explains why, for Blake, option c above (to blame Amy non-hypocritically), is not morally substitutable with option d above (to blame Amy hypocritically and initiate mutual deliberation between Amy and Ciara). The deeper level solidarity available to Blake and Amy, in virtue of their similar wrongs, is particular to *them*, and how they interact. Whilst Amy can recognise shared normative capacities with others, even other supremacists, the value of her solidarity *with Blake* cannot be achieved that way.

Other accounts of the value of non-hypocrisy have also emphasised the moral importance of acknowledging similar wrongdoing. However, it is worth noting some contrasts with the account given here. Matt King (2020), for instance, understands the problem of not acknowledging wrongdoing as a problem of mistaken moral priority. It is, King argues, more urgent for a person to attend to her own faults than those of someone else, and the fact that the hypocritical blamer attends to someone else's fault without attending to his own, shows a failure to appreciate that greater urgency.

One limitation in this account is that it does not tell us why *blaming* others makes a difference to the moral nature and gravity of not attending to our faults. It seems no less of a problem of mistaken priority, for example, to fail to attend to our moral faults when a similar fault grabs our attention in a TV programme. Mutuality, in contrast, can explain why acknowledgment has special significance in the context of blaming others. For the values of Mutuality depend precisely on the *combination* of blaming-and-acknowledging. Just as we cannot mutually engage others if we blame them and do not acknowledge our faults, we cannot mutually engage others if we simply acknowledge our faults and do not blame others. Obviously, we can have independent reasons to blame others, and independent reasons to acknowledge our wrongs; but only through doing both can we achieve the value of reasoning together.¹⁶

Now, Mutuality is not the only view that sees the combination of blame-and-acknowledgement as morally significant. A prominent alternative is the egalitarian view, such as Wallace's, that I described in Chapter 1.¹⁷ On these views, remember, when we blame someone and fail to acknowledge similar wrongdoing, we treat that person in a manner that we are not prepared to treat ourselves, and thus disregard their status as our moral equal. That is something we do not do, for example, when we *merely* fail to acknowledge our own flaws – here there may be defective self-treatment, but there is no unequal treatment.

The appeal to solidarity here does have an egalitarian flavour. However, it departs from the standard egalitarian thought in one central respect. Rather than the relational disvalue of treating a person as lower in moral status, my focus has been on the relational value of recognising ourselves and another person as appropriately included in the same normative

¹⁶ This offers a response to Gerald Lang's (forthcoming) sceptical remark: 'why', Lang asks, 'can't we detach issues regarding *self-application* from issues regarding *criticism of others*? Similarly, Dover (2019) might ask: Why "believe that non-critics are any less obliged to self-scrutinise than critics are"? (Dover 2019, n. 418)? The answer is simple: the value of mutual criticism requires the two things to be attached.

¹⁷ Again also see Fritz and Miller (2018), Lippert-Rasmussen (2013; 2018, ch. 4); Wallace (2010); Duff (2010, p. 127, pp. 138-139); Tadros (2009, pp. 394-404).

task. The emphasis isn't on showing ourselves not to be more morally worthy than others. Instead, it is on showing that we are involved in something together with others.

Conclusion

In sum, what is special about non-hypocritical blaming is that it allows blamer and blamee to address a certain range of subjects – their similar wrongs – multilaterally. I have also said that this process, in its ideal form, is more valuable than hypocritical blaming: mutually deliberative blaming improves prospects of corrective knowledge and motivation; it involves using the parties' agency in a fuller way to collaboratively reform; and it involves mutual recognition of a deep normative connection between agents who confront urgent moral reasons of the same type.

We are now in a position to see how hypocrisy, though good and apt, is also *inapt*. Notice how all of the values discussed here – corrective knowledge, collaboration and solidarity – are among those I listed in Chapter 2, in the account of the internal values of blame. Notice, furthermore, that a *version* of all of these values is possible in hypocritical blaming. We can have *some* collaboration, *some* solidarity, and *some* corrective knowledge even if a blamer blames hypocritically. Hypocritical blame is good and apt in that respect – it furthers blame's internal values. However, in each case, I have also argued that the version of these values in hypocritical blaming is only a second-best version. Non-hypocritical blaming enables better prospects of corrective knowledge, and richer collaboration and solidarity between blamer and blamee. This helps us to see the sense in which hypocritical blaming is also *inapt*: whilst it can instantiate blame's internal values, it does not instantiate them in their ideal form, or to their ideal extent. Hypocritical blaming is *inapt*, then, in the sense that it involves a *deficiency* in appropriateness, despite also being positively appropriate.

However, things that are non-ideal and *inapt* are not necessarily wrong. Let us now turn to examining why, in addition to being *inapt*, accurate hypocritical blaming is morally wrong.

Chapter 4. Non-Hypocrisy: Morally Required.

I have just argued that hypocrites *can* do better. But so what? It does not follow from the fact that a person can do better that they must. As noted in Chapter 2, we are not always morally required to bring maximal value into the world. So, to substantiate the view that hypocrisy is wrong, we need to explain why it is wrong for hypocrites to fail to advance the value of Mutuality in blaming. This chapter argues that what we need to do to enable that value – again, to openly acknowledge similar wrongdoing – is normally not costly enough to justify failing to enable it. We owe those we blame the low-cost benefits of acknowledgment, and so wrong them by not delivering those benefits.¹

The argument relies on a general duty to benefit others significantly at low cost to ourselves. I explain this duty in Section 1, and outline how I will apply it to non-hypocritical blaming.

In Section 2, I explore two expected costs of acknowledgment that might justify not doing it: psychological difficulty and effort. I argue that typically these costs are not high, and not as weighty as the benefits of acknowledging. I also point out that other principles of non-hypocrisy can be at least as demanding.

Beyond worries about demandingness, we might have deeper worries with the argument, two of which I consider in Section 3. First, that the blamee is not in a position to

¹ These benefits, I'll explain below, are not benefits in the most ordinary sense: they are opportunities for appropriate reflection on wrongdoing. But such opportunities are beneficial to a person in the sense outlined in Chapter 2 – people have reason to value responding appropriately to wrongdoing for their own sake, as part of a flourishing life.

claim additional benefits. Second, that the argument does not justify a duty of non-hypocrisy, but some other duty.

The anti-hypocrisy principle I argue for turns out to have a different form than the sort of principle we may have expected. The duty of non-hypocrisy is not a negative duty, which determines whether we ought to blame a person, but a positive duty, which tells us whether we ought to blame someone in one way rather than another. I defend this unorthodox idea in Section 4.

In concluding, I briefly consider how this analysis might illuminate instances where someone is not going to acknowledge wrongdoing when blaming (so is going to blame hypocritically if they blame), but wants to avoid acting wrongly, and thus decides to keep quiet rather than blaming the other person.

1. Low-cost beneficence

We have already seen an example where the low cost of providing benefits plausibly grounds a duty to provide it. Recall:

Parrot: You pose a very high risk to your life upon entering a burning building. Once you get to the only intact room, you see a parrot in a large cage, and a large baby. You can either flee with the parrot in the cage, or flee with both the baby and the parrot in the cage – but the latter makes for heavy and awkward lifting.

We observed that even though you are morally permitted not to enter the building, it is wrong to flee with just the parrot. Once you enter, you are required to save the baby along with the parrot. Here I will outline how similar considerations explain the wrongness of choosing hypocritical blame rather than non-hypocritical blame.

i. Claims to benefits

Consider again what explains why it is wrong to flee with only the parrot. Part of the explanation, I said in Chapter 1, must be to do with the superiority of available alternatives: saving the baby is morally superior to saving only the parrot. But that cannot be the whole

story, as saving the parrot is morally superior to doing nothing, and doing nothing is permissible. Furthermore, suppose that in order to save the baby, you would have to run through twenty more burning corridors to get to it – then, it is not so clear that you acted wrongly in fleeing only with the parrot, even if your moral reasons to save the baby are weightier than your moral reasons to save the parrot.²

So, another part of the explanation of what is going on in Parrot must be this: the further personal costs of saving the baby are outweighed by the moral considerations in favour of saving her. The cost of extra heavy lifting, that is, does not justify failing to save. This suggests that we have a duty, sometimes, to significantly benefit others where the personal cost of doing this is relatively insignificant.

I don't take this to be a controversial idea, though there is some disagreement about how best to formulate it. One view is that saving the baby is morally superior because this promotes more impersonal value in the world compared to saving the parrot. Worlds in which babies and parrots are rescued are morally better than those in which only parrots are rescued, and the goodness of that better world gives us a reason to create it which outweighs the significance of the fact that in it, you will have incurred the additional cost of heavy lifting.³ But we need not understand the duty to rescue in that particular way. Instead, we can think that rescue is owed to the baby. Your interest in avoiding heavy lifting is less weighty than the baby's interest in having her life saved. It is reasonable for the baby (or her representative) to reject a moral principle that would permit you to leave the building without her – in that sense, you cannot justify to her why you may not save her, and thus it is wrong not to save her.⁴

² Again, this draws on Harman's idea (2016) of a morally permissible moral mistake: considerations of personal cost can make it the case that you can be morally permitted not to act in the way that the balance of moral reasons recommends.

³ E.g. Theron Pummer's *Avoid Gratuitous Worseness* principle might be interpreted that way – see Pummer (2016), and for further discussion, see Sinclair (2018), p. 45; and PEA Soup (2017)

⁴ This is Scanlon's way of explaining (1998) why certain acts are wrong.

I defend a similar complaint on behalf of the blamee in hypocritical blaming: Amy's interests in Blake blaming her non-hypocritically rather than hypocritically outweighs Blake's interests in making the opposite choice. Granted, blaming is not really a case of rescuing, and the benefits of non-hypocritical blaming are less than the benefits of having one's life saved. Moreover, we do not have a duty to provide just any benefits to others at low cost; for example, it may be permissible not to give a well-nourished person a chocolate bar at the click of my fingers. But, as I argued in the previous chapter, the benefits of acknowledgement are morally significant. Briefly recall the content of these benefits: advancing our interlocutors' corrective knowledge and motivation; contributing more fully to the Collaborative Project; and building bonds of solidarity. These opportunities for appropriate correction, and reflection on wrongdoing, are not the same as chocolate bars. They might lead to greater or higher pleasure in the long term in a person who reforms, and a wrongdoer's reform might also make victims and others happier when this results in compensation or closure. However, the value of these things in a person's life does not derive solely from the pleasure they entail. Some of their value derives also from the fact that they are part of living well and flourishing as agents, and that they help us form certain bonds with other agents. These things may also make people happier, but as I have argued in Chapter 2, they can make a significant contribution in themselves to how well a person's life goes. This supports a strong presumption in favour of providing such benefits.

ii. Incoherent?

Parrot, remember, like some cases of hypocritical blaming, is an instance of the Worse-than-Nothing Problem, where it seems morally better to do the thing that flouts low-cost beneficence (e.g. save the parrot, or blame someone hypocritically) than to do nothing at all, and yet it is permissible to do nothing at all. We might doubt a duty of low-cost beneficence in *these* cases, even if we accept there is such a duty more generally, on the ground that it is incoherent. Can it really be the case that whilst you were permitted not to enter at all, you are forbidden from leaving the baby behind, despite having done something good?

Here is one way to articulate that worry. On the one hand, costs of entering the building justified not entering at all; but on the other hand, once you enter, if you take on further costs to save the baby too, your total cost intake will be higher. This means that if you are required to save the baby, you are required to bear *more* costs in total, for the sake of certain moral considerations, than the cost which made it permissible not to act, and hence not to respond to those considerations at all. How can that be?

But, this is to rule out what cases like Parrot strongly suggest, namely, that it is possible that the moral significance of costs can vary in the light of whether we have already borne, or decided to bear, these costs. I think that under certain conditions, the difference in cost between the costliest alternative (save baby and parrot) and cheapest alternative (not enter) can be a less weighty consideration in the morality of choosing between the costliest alternative (save both) and a slightly less costly alternative (save parrot). Once such condition is that we have already borne, or decided to bear, the costs of the slightly less costly thing.⁵ More than coherent, this is compelling in Parrot. You have *already decided* to enter the building and run risks to your life; so, appealing to *that* cost is not forceful in justifying a preference to avoid the total cost involved in the costliest, morally best alternative.

I will return to the Worse-than-Nothing Problem in the conclusion of this chapter. In what follows, I defend a similar claim about the morality of choosing between the alternatives of hypocritical and non-hypocritical blaming. The relevant costs in deciding whether we ought to blame non-hypocritically must be the additional costs of *acknowledging* wrongdoing, rather than the total cost of *blaming and acknowledging* wrongdoing; as the costs of blaming at all are ones which hypocritical blamers decide, by hypothesis, to bear. Those additional costs are the ones we need to consider when assessing whether we have a duty to blame non-hypocritically rather than hypocritically.

⁵ On the importance of having already absorbed costs, see McMahan (2018); and importance of a willingness to absorb costs, see Horton (2017).

iii. *Expected low-cost beneficence*

There is another difference, though, between the benefits of rescue in Parrot, and the benefits in non-hypocritical blaming. The former are guaranteed – I really can save the baby – whereas the latter, as noted in Chapter 3, are potential – a conversation takes two, and this is no less true of a mutually deliberative conversation. Blake cannot really foist the benefits of mutual deliberation on Amy; he cannot literally force her to discuss either her wrong or his wrong with him, at least if deliberation is going to benefit them properly. He can only hold out the olive branch, so to speak.

However, we can still think about the benefits of non-hypocritical blaming as *expected*. In general, I take it that if we ought to actually benefit others in some way, we ought, other things equal, to create a significant chance that they'll benefit in that way. Indeed, this rationale applies to blaming quite generally, at least on the account given in Chapter 2. Whether we blame in order to benefit victims, for example, or to benefit the blamee, these benefits will only fully materialise if the blamee is willing to absorb its content. Still, the fact that blaming provides a chance to reflect on wrongdoing is a core reason why hypocritical blaming is better than not blaming at all.

Furthermore, I submit that if we have duties of low-cost beneficence, we also have duties of expected low-cost beneficence. If, that is, we morally ought to benefit others at low cost to ourselves, we morally ought, other things equal, to do something that is expected to benefit them when the expected cost of doing this to ourselves is low.⁶

Let me turn to explaining why I think this duty applies to non-hypocritical blaming: in particular, why we can expect the costs of acknowledgment to be low.

⁶ This thought has been used to explain why risky behaviour in general can be morally wrong: see Parr & Slavny (2019). In its support, we could point out that the same relationship seems intuitive across a number of normative and evaluative claims. E.g. 'If my doing X is wrong, then it is wrong to increase the probability of my doing X'; 'If S is a bad state of affairs, then I have a reason not to increase the probability that S'; 'If my doing X makes me better off, then I have a reason to increase the probability that I do X'. And so on.

2. Acknowledgment: cost or benefit?

Before proceeding, I should say that nothing below denies that the costs of acknowledging wrongdoer matter, and can even be decisive. Sometimes these costs will have *exempting* or *excusing* force: the thought of acknowledging our wrongs may induce panic, for example, making us psychologically unable to go through with it; and this might modify or negate our culpability for blaming hypocritically instead. Sometimes costs will have *justificatory* force too. Depending on the nature of the wrong in question, and the criminal and social status of that wrong, there will be cases where acknowledging it will lead to serious harm, such as trauma, social stigmatisation, imprisonment, or violence against us – where that is true, the failure to acknowledge may be permissible (though it is not clear that such a failure can be appropriately called hypocritical). Lastly, there will be cases where we have benefitted others so much in the past that the moral pressure to benefit *this* person runs out of steam. But, there are excusing and exempting conditions, and cost-thresholds, on most duties, so this is no special problem for the duty of non-hypocrisy.

The two costs of acknowledgment I will focus on are less extreme. They are what I take people to be typically worried about in ordinary blaming circumstances: namely that acknowledgment will be shameful, and that it will be laborious. Neither, I argue, is a good reason to blame hypocritically.

i. Shame

When outlining the Worse-than-Nothing Problem in Chapter 1, I considered whether accurate hypocritical blaming might be worse than no blame because it is unpleasant – it is not nice to be publicly confronted with our faults, and it is not nice to know that others think badly about us (indeed it might be bad for us that others think badly of us full stop). Similarly, when discussing overblaming in Chapter 3, I suggested that we ought ideally to ensure that whilst strict in tone, to reflect and communicate the gravity of wrongdoing, blame should not impose more costs than are necessary for a person to achieve these aims. But obviously, in the real

world, blame barely ever meets the ideal: a blamee can feel ashamed even if blame's content or tone does not warrant that feeling, and even where the blamer has no intention to cause such feelings.

Now, if these are valid concerns about blame, they should also be valid concerns about acknowledgment. Acknowledgment is about being willing to reflect on our wrongs with another person; it is a way of publicly confronting our responsibility. If that can be shameful for the blamee, we should allow that it can also be shameful for the blamer. This might lead us to suspect an important difference in the nature of costs between Parrot and Supremacists. In Parrot, the heavy lifting needed to save the baby may be costly, but that kind of cost does not seem quite as stark as the cost of publicly examining our flaws – heavy lifting is a drag, but it is not psychologically piercing. So the difference in costs between the best and second-best options in Parrot is perhaps less serious compared to the difference in costs between hypocritical and non-hypocritical blaming.

However, whilst severe shame in acknowledging wrongdoing can be an important factor, it is unlikely to be decisive in many cases.

a. Acknowledgment and Revelation

One reason to question the idea that shame normally permits us not to acknowledge is that acknowledgment has a limited effect on shame when our faults are already revealed to those we blame. To see the difference between revealing and acknowledging, consider an interaction between Amy and Blake that is analogous to a case I gave in Chapter 1:

Shamelessness 2:

Blake: That was wrong of you to make that racist remark.

Amy: Fine, but same for you; I witnessed your casual casteism too!

Blake: You really must think about what you did, and apologise.

Amy: But you *saw* me witness your casteist remark!

Blake: No chance.

Amy: Sigh.

Typically, we feel ashamed by our flaws when and because they become known to others. But, as the above example shows, there are many cases where this cannot be a source of shame whilst blaming others; for the blamee knows about what we did, and we know that they know this. Any shame Blake might reasonably feel in this case is due to the further prospect of *reasoning about* her flaws with others. This is less shameful, I think, than the prospect of revelation.

Of course, not all hypocritical interactions are as above: blame is not always like Shamelessness 2, or indeed the Sage Supremacist case that I described in Chapter 3, where similar wrongdoing is mutually known. However, this is not an artificial stipulation; it is actually pretty mundane. Very often this is *why* we get riled up about hypocrisy in the first place – we get riled up *because* we know that the blamer is guilty of certain faults, and that they too are aware that we know this. The problem is that these people refuse to reflect on their faults with us. Granted, mutual reflection can also be shameful. Discussing our wrongs involves lengthy public focus on them, and that lengthy public focus may be humiliating. Still, the point is that the shame involved in such cases is less of a problem than in cases where others do not know about our wrongs, and thus the discussion about what we did requires us to reveal them.

b. Shame mitigated

What, then, should we make of the latter cases? Suppose Blake has to both reveal his wrongdoing and publicly scrutinise it in order to enable mutual deliberation with Amy. Even here, I think the significance and likelihood of shame can be overstated. We have to remember that one subject is already out in the open; namely *Amy's* racism. Given this fact, *Amy* knowing about Blake's casteism is not an especially shameful prospect. Shame in revealing our wrongs is most acute when we are concerned that others will see us as morally inferior, or our behaviour as especially or strikingly morally base. But that sort of concern will be minimal when we are already discussing the other's similar wrong. Blake, for example, should not worry that Amy will consider him inferior, or his casteism especially or strikingly morally base, when

her own racism is already under the spotlight. This heavily mitigates the intensity and likelihood of his shame.⁷

c. Reaction costs

Furthermore, in considering the emotional costs of non-hypocritical blaming, note that *hypocritical* blaming can also have emotional costs for the blamer. One such cost is a negative reaction by the blamee. In deciding whether and how to blame Amy, for example, Blake might legitimately fear Amy lashing out at him defensively; or he might fear making her feel overly guilty or ashamed. However, such costs are greatly reduced in non-hypocritical blaming, for a reason already given above: Blake's acknowledgment shows that he does not mean to present himself as morally superior, or to portray Amy as strikingly morally base. This will help preempt Amy's defensive anger, as well as her sense of shame.

So, whilst acknowledgment might increase emotional costs in some respects – shame – it decreases emotional costs in others – reaction costs. Hence a further reason to deny the general significance of these costs in grounding a permission to blame hypocritically rather than non-hypocritically.

ii. Effort

I noted above how acknowledgment and revelation come apart: what we're after is not merely that others *know* about our similar wrongs; these wrongs must also be available for common scrutiny. This point brings to light a different burden that Mutuality might impose. Beyond risking emotional costs, a process of common scrutiny – of listening and responding to others' views about our wrongs – is likely to take some effort. However, I doubt this effort is normally too demanding an ask.

⁷ Shame will remain a real concern when third parties, whose faults are not under discussion, witness the blaming interaction. But it may still be considerably mitigated, as shame about arriving late to meetings is mitigated when walking in with a fellow latecomer.

a. Not always taxing

One thing to note is that although mutual deliberation does involve a committed effort, acknowledging wrongdoing does not always mean undertaking it. I have pointed out that it takes two to tango. Blake and Amy cannot achieve the value of Mutuality if any one of them is unwilling to play their part, and where Amy is unwilling, no mutual deliberation process will occur. The duty to acknowledge does not, as it were, require us to converse with ‘a brick wall’. (Again, it does not seem hypocritical to ‘abandon ship’ once we can see that mutual conversation is not on the cards.)

Furthermore, where acknowledgment does require the effort of mutual reflection, it will not always be as onerous as we might expect. As Dover (2019) points out, criticism need not be a ‘drive-by’ event, where critics rifle moral considerations towards one another, wait for appropriate responses, and then swiftly move on (2019, pp. 400-402). People can sustain a conversation for as long as they wish; and, they might sustain one even when attending to other things, by pressing pause, and picking up where things left off at a later point. I have argued that wrongdoing is a morally pressing matter, and we should respond to it urgently. But speedy reform is often not realistic, as careful reflection on wrongdoing can require time and space. The process of mutual deliberation begins to look less taxing when the effort is spread over an extended time-period.

b. Other principles

Another thing to note is that Mutuality is less demanding than more orthodox views of wrongful hypocritical blaming. Compared to the principle defended here, alternative, popular principles of non-hypocrisy imply that at least as much effort is needed to avoid blaming wrongly.

One common way to understand the duty to blame non-hypocritically, hinted at in Chapter 1, is that it is hypocritical, and therefore wrong, to blame others for similar wrongs

full stop, whether we acknowledge wrongdoing or not. We can express this view as the following principle:

Blame innocently: You may blame others only if you are innocent of similar wrongdoing.⁸

On this principle, blaming violates moral constraints whenever the blamer is guilty of similar wrongs. The only way to act permissibly is thus either not to blame others at all, or, if we are going to blame others, not to commit the type of wrongs we blame others for. Thus, given that Blake has committed casteism, his blaming Amy will be at least in one way wrong, and there is nothing he can do to change that fact.

A second view about wrongful hypocritical blaming is that it is hypocritical, and therefore wrong, to blame others when our commitment to moral values is lacking. We can express this as follows:

Blame committedly: You may blame others only if you are committed to the moral values flouted by the blamee.⁹

On this principle, blaming violates moral constraints whenever the blamer is not committed to the values the blamee has flouted. The only way to act permissibly, then, is either not to blame people guilty of flouting these values, or to be committed to those values. So long as Blake lacks commitment to values that are flouted by prejudicial abuse, for example, his blaming Amy is at least in one way wrong, and there is nothing he can do to change that fact.

I take it that both principles make it very difficult to avoid acting wrongly. Our lives are full of moral imperfections: we make mistakes all the time, and our commitment to moral values tends to be weak or fluctuating. Moreover, many of us are prone to criticising others in our interactions. Avoiding wrongdoing, building and sustaining moral commitments, and

⁸ I think Gerald Lang (‘forthcoming’), Daniela Dover (2019) and perhaps Macalester Bell (2012) tend to interpret the duty this way.

⁹ Patrick Todd (2017) defends a cousin of this view, namely that uncommitted blamers lack the moral right to blame (which he equates with lacking standing to blame). The principles differ to the extent that we can permissibly do what we have no right to do, and can have a right to do wrong. I explore Todd’s actual view in more detail in Chapter 5. See Rossi for development of this view of hypocrisy (2018).

refraining from criticism if we fail in these others areas; all of these things are not likely to come easily. Thus, it will take a great deal of effort not to violate non-hypocrisy.

Now compare those alternative views with the principle on offer here:

Blame concessively: If you blame, or decide to blame others, you ought to be ready to mutually deliberate about your own similar wrongdoing with those others.

On this principle, we may act permissibly by either not blaming others at all, or by blaming others and acknowledging wrongdoing. The fact that Blake blames Amy does not necessarily mean he has acted wrongly by blaming her, as when he does blame her, there is still something he can do to avoid acting wrongly: acknowledge wrongdoing. Though acknowledgment will likely involve some effort, that effort is much lower than the effort that is needed to avoid wrongdoing altogether, or the effort that is typically needed to build moral commitment to certain values, plus the effort needed, if we don't do these things, to ensure that we refrain from criticism.

I will come back to these alternative principles later. Right now, the point is just that if we do not find it implausible that morality would demand these alternative principles, we should not reject the duty to acknowledge simply because it takes effort.

iii. Moral benefits

That, of course, is no particular reason in favour of Blame Concessively – it might be that all three principles are too demanding. But in judging whether the duty to acknowledge is too demanding, we need to expand the analysis so far. So far, we have evaluated costs of hypocritical versus non-hypocritical blaming from the standpoint of someone who prefers not to engage in mutual deliberation on the ground that it is laborious or shameful. Sometimes these preferences will be unmet. However, we have not yet compared the costs from the standpoint of someone who wants to flourish as a moral agent. Once we do, it becomes clear that acknowledging is often beneficial for the blamer overall, rather than costly overall.

So far, I have argued that Blake ought to enable mutual deliberation to give *Amy* a chance to benefit morally from it: through improved correction, collaboration, and solidarity.

But remember that Blake is a similarly morally situated agent, whose own flourishing also involves appropriate responses to his wrongdoing, engagement in collaborative exercises of agency, and forming solidaristic bonds with others. In other words, the benefits of mutually deliberative blame are mutual. If Amy has interests in these moral benefits, so does Blake. He does not get these benefits from the interaction where he blames hypocritically. That makes non-hypocritical blaming in one significant way more beneficial for Blake.

As mentioned earlier, other costs can sometimes outweigh these moral benefits of acknowledgment. Self-critique may sometimes be seriously crushing, and thus worse overall for the target, even if better from a moral point of view. But *normally*, our interests in being blamed outweigh our interests in avoiding blame – this was a central driver of the Worse-than-Nothing Problem in the first place.¹⁰ A key source of the plausibility of the idea that accurate hypocritical blaming can be better than no blame is that accurate blaming, even if unpleasant, is morally worthwhile. The same holds true for accurate acknowledgment. Where Blake’s interests in appropriate correction are more significant than the interests in avoiding effort and emotional discomfort, non-hypocritical blaming is better for him overall. This, I suspect, will typically be true.¹¹

¹⁰ This is also why Wallace’s argument (2010) struggles with the Worse-than-Nothing Problem – it overlooks the fact that blame can be beneficial for a person overall even if somewhat harmful. Here is another way of stating the objection raised in Chapter 1: Suppose I do something that harms someone a bit, yet benefits them overall, say by giving them some awful medicine to treat an illness. It is hardly an objection that I treat this person *unequally* if I am similarly ill and don’t give myself the medicine. I think blaming can be analogous, and thus not a serious form of unequal treatment.

¹¹ Tadros (2016; 2020b) likewise argues that the interest in not being responsible for wrongdoing can matter a great deal more than other significant interests, by appealing to what parents should prefer for their children. For example, a parent intuitively has reason to prefer that their child suffers a broken leg rather than wrongly breaking another child’s leg (Tadros, 2016). Similarly, a parent intuitively has reason to prefer that their child is wrongly accused and punished for wrongdoing rather than being responsible for wrongdoing (Tadros, 2020b). We might read the present discussion as extending that idea to

And non-hypocritical blaming has a further moral benefit for Blake. A reason I gave in Chapter 1 for why not blaming others at all can be permissible is that blaming someone threatens to disrupt morally valuable relationships we have with them, by fuelling anger and resentment. But, again, that threat reduces when we blame non-hypocritically rather than hypocritically. Amy distancing herself from Blake, or thinking that Blake is distancing himself from her, for reasons indicated earlier, is far less likely when Blake acknowledges similar wrongdoing. In that case, Amy is less liable to resent or shield herself from his critical attitude, seeing it as less of a personal attack; and she is less liable to perceive him as distancing himself from her, given his public recognition that they are in the same moral boat.

Two clarifications before proceeding to different concerns about the duty to acknowledge. First, it is often the case that considerations of personal moral benefit are less morally urgent than considerations of benefit to others. Perhaps I act more wrongly in failing to help someone else compared to when I fail to help myself. But recall that the duty to blame non-hypocritically is not grounded in personal benefit alone – it is grounded in the low personal cost of benefitting others. The fact that acknowledgment furthers my own interests in living well makes it, in one way, less costly, and thus makes it harder to justify blaming hypocritically instead.

Second, the argument here is not that Blake morally benefits from acknowledgment *because* it is the right thing to do. That would beg the question – it would merely assume that we ought to blame non-hypocritically rather than hypocritically. Rather the explanation goes the other way: part of what makes acknowledgment the right option is that Blake stands to benefit morally from doing so. This reduces the cost to him of benefitting Amy morally, and thus helps explain why he owes her a duty of low-cost beneficence.¹²

responsibility for appropriately correcting for our mistakes: parents have reason to prefer that their children are responsible for correcting their wrongs at the cost of a broken leg, let's say, rather than that their child avoids a broken leg and is not responsible for correcting their wrongs. This further suggests that we should prefer correction for our own sake at significant cost.

¹² Thanks to Simon Gansinger for discussion.

3. Further worries

A different way of criticising the duty to acknowledge is to agree on what I have claimed so far – to agree that acknowledgment benefits others at low-cost – but to challenge the conclusion – that we therefore have a duty to blame non-hypocritically. This section considers two such worries.

i. Good blaming means good non-hypocritical blaming

First, some might deny that the blamee is owed any benefits at all. Unlike the beneficiary in Parrot – the baby – the person we can benefit by accurate non-hypocritical blaming is not an innocent party, but a blameworthy wrongdoer. Do such people really have a claim to benefits? Indeed, we tend to think that wrongdoers are often not wronged when people actively *harm* them, rather than merely fail to benefit them, such as when we impose suffering on criminal wrongdoers, along with severe limits on their freedom, on the grounds that they make themselves liable to bear these costs. Perhaps it seems odd to think that such harms can be required at the same time as us owing wrongdoers benefits like those of non-hypocritical blaming. However, I doubt anyone seriously thinks that wrongdoers lose all entitlements, and many think that they have entitlements that are intimately connected to harms we impose on them. For example, we typically aim, in criminal trials, to explain to wrongdoers why they are liable to the relevant harms. We also aim to allow them a chance to answer to accusations and sentencing verdicts, even when we know they are guilty of accusations. In other words, harms to wrongdoers are consistent with claims to certain accountability procedures.¹³ Moreover, the idea that wrongdoers hold at least some entitlements is already implicit in one of the intuitions that this thesis aims to explain – that hypocrisy is an interpersonal wrong, and the primary complainant in hypocritical blame is the blamee. That idea alone strongly suggests that wrongdoers' interests matter.

¹³ This is a major theme in Duff (2018a).

Still it might seem that wrongdoers' claims ultimately have nothing to do with wrongdoers: the reason we hold wrongdoers accountable cannot be for their sake; it is for the sake of victims, and other stakeholders in what they have done, such as, families of victims and people with interests in not being victims of similar wrongdoing in future.

Whilst this is a possible view, considerations offered in Chapter 2 speak against it. There are cases of warranted blaming where no accountability process is owed to victims, as when we harm a non-human animal, for instance, or when we wrongly set fire to one of a few remaining copies of an ancient book. Engaging in accountability in such cases is still apt and good for the blamee, and can provide a reason to blame them for their own sake. This is what I have suggested is the core point of blaming. What sets it apart from other kinds of response to wrongdoing is the need to engage with wrongdoers about their wrongdoing, and that is centrally for their sake, irrespective of whether others benefit. That is why it is better than not blaming in the first place.

Some will no doubt take the opposite line on grounds of desert. One view is that harm, rather than benefit, makes blaming good – deserved harms provide reasons to blame rather than keeping quiet.¹⁴ The account of blame I gave in Chapter 2 did not reject this view. But it poses no real challenge to what I have said. Where Blake accurately and proportionately blames Amy, he is (by stipulation) already going to subject her to whatever harm is meant to be deserved. Now, *once* people get what they deserve, benefits to them can surely matter to how we ought to treat them; I doubt there is any good argument for why, when people have already been harmed as they deserve to be, they have no claim against further harm, or no claim to any benefits. So it is plausible, even on a desert-based view, that benefits to the blamee of non-hypocritical blaming are morally important.

¹⁴ E.g. Nelkin (2016); McKenna (2012a; 2012b). Even this view might ground a reason to acknowledge in a different way. A possible argument is that Blake ought to allow himself and Amy to mutually harm one another through reasoning together, in order to fully deliver just deserts to each other. Of course, this is not my view.

I also think the idea that wrongdoers don't have claims that we benefit them becomes unconvincing when we consider again the nature of the benefits under discussion. The importance of appropriate correction is not that it is pleasurable; indeed achieving correction often requires participation in a painful process. So, we are not rewarding wrongdoers for bad behaviour; we are merely urging them to correct for their bad behaviour in a certain way, even in a costly way. Desert-based views cannot plausibly condemn this way of benefitting blamers, as that, absurdly, would mean condemning most cases of successful accurate blaming on grounds of desert.

ii. *Non-hypocrisy as discursive willingness*

A final objection to the argument in this chapter is that even if we do have a duty to do what the argument suggests, this is not a duty of *non-hypocrisy*. For the type of acknowledgment required by blamers on this account – a willingness to reason about our wrongs with those we blame – is not needed to avoid hypocrisy. One reason why some might think this is because they see hypocrisy as a purely internal matter; a matter of how people privately view themselves relative to others, and what people care about deep down. Privately, Blake can fully acknowledge his casteism without discussing it with Amy: he can appreciate and evaluate his faults in his head; he can be angry or upset with himself; and he can be morally committed against repeating the offence.¹⁵ So long as he does this, he is not hypocritical at all, even if he is not ready to publicly acknowledge. Another reason to question whether mutual deliberation is a duty of non-hypocrisy is to hold that the crucial question is whether Blake *admits* what he did.¹⁶ To avoid hypocrisy, some might claim, it is not necessary that Blake be willing to reason about his casteism. It is sufficient that he admits to his casteism – once he does this, he no

¹⁵ For the view that hypocrisy is an internal matter, see e.g. Fritz and Miller's (2018) and Todd's (2017) criticisms of Wallace (2010). See also Duff (2010, p. 128). Chapter 5 provides further discussion on how to assess private hypocritical blaming.

¹⁶ Thanks to Stephen White for raising this point.

longer blames hypocritically. Thus, some might conclude that the duty of Mutuality – to be willing to publicly reason with others – even if there is one, is not a duty that is distinctive of hypocrisy.

It is plausible that blaming is less hypocritical where we privately acknowledge wrongdoing compared to where we do not. But it is also plausible that open acknowledgment – especially acknowledgment which shows a willingness to reason with others – makes a difference to the hypocrisy involved. This came out in another conversational example I gave when discussing the Moral Makeup view in Chapter 1:

Diversion 2:

Blake: That was wrong of you to make that racist remark.

Amy: Fine, but same for you; I witnessed your casual casteism too!

Blake: Sure, but you really must think about what you did, and apologise.

Amy: But you *saw* me witness your casteist remark.

Blake: Stop changing the subject!

I take it that Blake behaves both hypocritically and wrongly here, even if he sincerely admits his casteism. If he is sincere, then he privately acknowledges his casteism. Yet that, clearly, is not enough to escape a complaint of hypocrisy.¹⁷ Thus, there is an intuitively significant contrast between hypocrisy and the willingness to openly reason about similar wrongdoing when blaming others. And this, I think, is all we need to vindicate a duty of non-hypocritical blaming: to show how at least one type of hypocrisy involved in blaming can be wrongmaking.

But not only that. The wrongness of failing to mutually deliberate with others can also explain the importance of private acknowledgment of wrongdoing, in the following way. Publicly evaluating our faults is part of the ideal way to blame similar wrongdoers, as only by doing this can wrongdoers reason together about their similar faults. But, if it is important that we publicly self-evaluate, it is also important that we privately self-evaluate. Unless we are ready

¹⁷ For slightly different views, compare Roadevin (2018) on reciprocity (Roadevin, 2018, pp. 142-145), and Rossi (2018) on the disposition to accept blame (Rossi, 2018, pp. 562-563)).

to evaluate our behaviour in our own minds, we cannot be ready to evaluate our behaviour in conversation with others. So, the importance of public acknowledgment explains the importance of private acknowledgment, as the latter is a necessary part of the former.

The above also illuminates the role that deception might play in a complete account of wrongful hypocrisy. At least part of the duty not to deceive others about our wrongs when blaming them derives from the duty of willingness to deliberate mutually about those wrongs. As I explained in the previous chapter, acknowledgment means putting our wrongs ‘on the table’ for discussion, and wrongs that we successfully deceive others about are not mutually known, and thus are ‘off the table’ for discussion; we cannot appropriately discuss things we are unaware of. Thus, we ought sometimes to avoid deceiving others for the sake of deliberating mutually with them. In cases like *Shamelessness 2*, where we know that the blamee knows about our similar wrongs, the duty to mutually deliberate is not violated by what Lackey (2013) calls *deception*; but, it is violated by what she calls *being deceptive*.¹⁸ In *Shamelessness 2*, even though Amy knows about his casteism, Blake still play-acts in his outright denial of it, and that sort of play-acting is clearly an impediment to mutual discussion of the topic. So, both being deceptive, and deception, can be prohibited on grounds of hypocrisy.

On the other hand, our duties of non-hypocrisy do not derive *solely* from our duties to avoid deception or being deceptive, as *Diversion 2* illustrates. In *Diversion 2*, Blake is neither deceiving Amy nor being deceptive. Rather than play-acting, he admits to his wrongdoing. Still, he acts both wrongly and hypocritically, as his diversion is equally an impediment to mutual discussion. This explains both why the factor of deception matters, and why it does not quite get to the heart of the matter.

4. Whether versus how to blame

Although the type of hypocrisy I have argued against is recognisable, the duty it grounds is somewhat less so. But this, I think, turns out to be an advantage, rather than a problem.

¹⁸ Lackey (2013, *passim*) – see Chapter 1, sec. 1.ii.

Ordinarily, duties of non-hypocrisy are thought to be *negative*: we tend to see hypocrisy as something that needs to be avoided, and thus to assume that the wrongness of hypocrisy lies in doing something we ought not to do. In contrast, the duty of beneficence is *positive*. Rather than requiring Blake *not* to do something, it requires him to do something – to acknowledge and enable mutual deliberation if he blames. Thus, wrongful hypocrisy can be avoided by doing something, rather than by not doing something.

The contrast comes out when comparing the non-hypocrisy principles I earlier outlined. Again, compare the three alternatives below:

Blame innocently: You may blame others only if you are innocent of similar wrongdoing.

Blame committedly: You may blame others only if you are committed to the moral values flouted by the blamee.

Blame concessively: If you blame, or decide to blame others, you ought to be ready to mutually deliberate about your own similar wrongdoing with those of others.

The first two principles are negative, and it is clear what they tell us not to do, namely, *blame*. Those principles say, ‘under such-and-such circumstances, do not blame!’ The third principle, on the other hand, has a different structure. It says, ‘under such-and-such conditions, acknowledge!’ It tells us what positive course of action we ought to take if we decide to blame.

In philosophical literature, hypocrisy is also commonly construed as a negative constraint. Here, for example, is how Dover (2019) describes the duty of non-hypocrisy:

By objecting to [...] hypo-criticism, [we] typically [imply] that [criticism] *ought not* to have happened, in light of the fact that [the would-be criticise made a similar mistake] (2019, 406, my emphasis).

And here, similarly, is how Kyle Fritz and Daniel Miller (2018) articulate their normative view:

[...] the fact that an instance of blame is hypocritical with respect to a violation of some norm N is a moral reason that *counts against blaming* in that case (2018, p. 506 my emphasis).

On the negative view, the significance of hypocrisy is that it bears on our decision *whether* to blame. The problem of hypocrisy is a problem in the act of blaming itself. So hypocrisy considerations play a role in guiding our decision whether to blame a person at all.

However, that view unnecessarily restricts the possible ways of objecting to hypocrisy. We might instead want to say that blame or criticism ought to happen, but if it does, it ought

to happen in a particular way: non-hypocritically, rather than hypocritically. On this view, hypocrisy considerations guides us in our decisions about *how* to blame – the problem is not about the blaming or criticism itself, but instead about its manner and ingredients. This is in fact a credible view, for our assessment of wrongful hypocrisy, as we’ve seen throughout, is clearly sensitive to what happens in the course of blaming: in particular, to whether the blamer acknowledges wrongdoing.

Why go for the negative view? One reason might be that we tend to get particularly angry about hypocrisy; we tend to see it as a serious wrong. And often, things we do seem morally worse than things we fail to do; for example, it is worse to put a baby’s life at risk by placing her in a burning building than to fail to take that risk away by rescuing her. But, failures to act are not morally trivial: my failing to save the baby is clearly a serious problem! And the distinction between acting and failing to act is not so clean; often we actively do things make it the case that we fail to do things.¹⁹ You can actively form an intention not to save the baby, and I can actively divert discussion away from my wrongdoing to avoid pressure to acknowledge it.

A second reason to characterise non-hypocrisy as a negative duty is that this captures how people typically respond to hypocritical blaming. We often say things like, ‘Hey, who are you to blame’, ‘how dare you!’, and ‘*you* can’t talk!’ We don’t say, ‘Hey, you are blaming me in the wrong way!’, or ‘hey, there is a better way for you to do this!’ But in this instance, I do not think our language reflects what we ultimately care about. When those blaming us make sure to place their similar faults on the discursive table, the relevant ‘how dare you’ feelings tend either to be much less intense, or to disappear entirely. Where frustration does remain about a fully acknowledging blamer, it is difficult to see what would justify this. Blaming non-hypocritically, after all, is blaming in the morally best way. Our blaming norms should encourage best practice, rather than no accurate blaming at all.

19. On this well-noted point about the moral difference between doing and allowing, see e.g. Smith (2007).

This raises a further challenge for negative duty accounts. Since doing nothing does not violate moral constraints, those who posit a duty not to blame face the difficulty of explaining why we have stronger moral reason to violate constraints than to keep quiet, and thereby violate no constraints. I raised this challenge for the egalitarian view in Chapter 1, which holds that hypocritical blaming violates a duty not to treat others unequally. But it will apply to other negative duty accounts. In contrast, a positive duty account faces no pressure to deny that hypocritical blaming is a moral improvement on silence, since it does not tell us that we have a moral duty to avoid criticism.

Conclusion: Discouraging second-best?

I have argued that the duty to blame non-hypocritically rather than hypocritically is explained, not by the badness of blaming, but the goodness of providing the blamer, and ourselves, with the morally best form of blame. As we ourselves stand to benefit from mutual deliberation in various ways, this benefit is not normally costly to provide to others. This helps us to see why a certain kind of hypocrisy – the unwillingness to reflect on our similar wrongs with those we blame – is morally wrong.

One final issue to consider relates to the Worse-than-Nothing Problem. I have said that the right account of hypocrisy should recommend the morally best blaming practices, rather than discourage blaming at all. But we should not forget about second-best blaming practices: surely the second-best option of accurate hypocritical blaming, should at least sometimes be encouraged – especially when we know that best option of non-hypocritical blaming will not be pursued.

But does Mutuality struggle to explain why? If, as I have argued, accurate hypocritical blaming is wrong, it might seem to follow that certain people will be discouraged from the second-best option of hypocritical blaming. Consider someone who does not want to act wrongly, is willing to blame hypocritically, but is unwilling to blame non-hypocritically –

unwilling, that is, to acknowledge wrongdoing.²⁰ Such a person might reason: ‘I am willing to blame this person hypocritically, but that would be wrong, and not blaming them is permissible. So, I’d better keep shtum.’ It seems problematic that morality, in making the second-best option wrong, would in this way discourage a person from doing something that is morally better than nothing.

This issue raises a puzzle of much wider scope than that of this thesis. But I think there may be a way to interpret what I have said without it implying that second-best conduct is discouraged. We can understand the wrong of hypocritical blaming as *contrastive*: its wrongmaking feature derives, not from the fact that it is chosen, but from the fact that it is chosen *rather than* a certain alternative: namely, the best alternative, of blame-plus-acknowledgment.²¹ Where an agent has decided not to pursue that best option, as is true of the person who will refuse to acknowledge, this wrongmaking feature of choosing the second-best option should figure differently in her deliberation about what to do. In that case she should contrast the wrongful, hypocritical option, not with the off-the-table wrongmaking options, but instead with the remaining options on the table, and act accordingly. Thus, the person who obstinately will not acknowledge should blame hypocritically, provided this is accurate and proportionate.

Here’s a tentative explanation for why that might be true. It might be that, to permit reasoning *simply* from the fact that something is wrong to the conclusion that I should choose a permissible alternative would be to permit ‘fetishising’ wrongness over what explains it. We might think morality cannot appropriately fetishise wrongness in that way. The non-acknowledging person should not do that either: he should not fetishise the fact that hypocrisy is wrong over what explains this fact; and what explains this fact, I have claimed, is not the contrast between blaming and not blaming, but the contrast between hypocritical and non-hypocritical blaming, the latter of which is now off the table.

²⁰ Thanks to Joe Horton for discussion; for his own solution, see Horton (2017).

²¹ For the view that normative reasons generally have a contrastive structure, compare Snedgar (2017).

Whether this works more generally is an issue for another project. In the remainder of this thesis, I will consider how the duty of Mutuality sheds light on a more specific moral issue: of why hypocrites ‘lack standing’ to blame others hypocritically, and what hypocrisy’s relationship with standing reveals about its moral significance.

Chapter 5. Standing and Procedural Value.

So far, I have explained why I think hypocritical blaming is wrong: non-hypocritical blaming advances the values of Mutuality – values that hypocrites ought to enable at the low cost of acknowledging wrongdoing. But that does not give us a complete picture of the moral significance of hypocrisy. We don't simply judge that hypocritical blaming is morally wrong; we tend to discuss hypocrisy in connection with a broader set of blaming faults of a certain kind. Philosophers use a specific term to describe this fault. The hypocrite, it is alleged, *lacks standing* to blame. However, there is no real consensus among philosophers on what this means in moral terms, and it is not clear what to make of it. Once we can see that hypocrisy is wrong, what do we gain by introducing a technical term – 'standing' – to describe this wrong? Furthermore, why do we use the same word to describe other instances of defective blaming? The rest of this thesis considers these issues. Investigating the relationship between hypocrisy and standing will help us see both what hypocrisy has in common with other standing-related faults, as well as its distinctive impact on standing. It will also clarify what role the hypocrite's loss of standing plays in explaining why her blame is morally wrong.

The idea that the blamer has defective standing to blame comes out in a certain kind of instinctive response to being blamed. I have described some of these already, but here is a longer list of examples:

'You can't talk to me about that!'

'You don't have a leg to *stand* on!'

'You don't *get* to blame me!'

'You have *no right* to complain!'

‘You’re in *no position*...!’

Such responses often target hypocritical blaming: as I mentioned at the outset, it seems apt for Amy to object to Blake’s blame in this sort of way. But the responses are equally compelling in other cases. Two examples, which I will explore in this chapter, are when people blame us for things that are ‘none of their business’, and when we learn that the blamer entrapped the blamee into acting wrongly – it is natural to object to the stranger and the entrapper in a similar way.

But notice a peculiar, *ad hominem* feature of this way of objecting to blame. What the ‘you’ responses above take issue with is not the content of blame, or its accuracy, or the appropriateness of its tone or force. Rather they draw attention to something about *the blamer*, alleging that certain facts about the blamer make their blaming objectionable. How can this allegation work? Isn’t blaming just an opportunity to focus on our responsibility for wrongdoing? Provided the blamer does this accurately and proportionately – provided the content, tone and force are right – it is unclear how personal facts about the blamer could make the blame objectionable. And, even if the blame *is* objectionable in some way, why should we care? In what way would personal facts make a moral difference to whether the blamee should be urged to see the moral truth?

In the next chapter I consider the second challenge – in what sense we should care about standing, and what sort of moral property it is. Here I consider the first challenge: I defend a particular view about why certain facts can (and do) affect a person’s standing. Certain facts about a blamer affect her standing, I argue, because these facts affect the value of the accountability procedure she engages in. I call this view *Procedural Value*.

In Section 1, I present this idea, and show how it explains the hypocrites’ defective standing to blame. Procedural Value also helps us understand the ‘none of your business’ and entrapment cases, in ways I describe in Section 2.

To gauge the success of this account, it helps to look at others. In the second part of the chapter (Sections 3-6), I consider the views that hypocrites (and others) lose standing because of their moral ledger, or because of their bad character, or because they violate equal

treatment norms. I will argue Procedural Value has much wider explanatory appeal than these alternatives. It better explains a whole host of relevant factors, such as entrapment, estrangement, and the impact of acknowledging wrongdoing. Procedural Value also provides a more compelling response to scepticism about the significance of standing.

1. The hypocrisy case

For non-sceptics about standing, hypocrisy is among the clearest cases where standing seems threatened. The Mutuality account can help us begin to see why.

Recall the explanation I gave for why it is morally wrong for Blake to blame Amy hypocritically for racism, even if this is better than nothing. Even where Blake has stronger moral reason to blame hypocritically than not at all, he has even stronger moral reason to blame non-hypocritically, and he cannot justify not complying with that strongest set of reasons in favour of the second-strongest set. This view included a parallel thought about what grounded this hierarchy of reasons. The thought was that accurate hypocritical blaming advances certain values – the values of correction, collaboration and solidarity – which ground reasons to blame hypocritically rather than to keep quiet. However, these same values are not as well-advanced as they might be: they'd be better advanced if Blake instead blames Amy non-hypocritically, with a readiness to reason with her about his casteism. That is why he has even stronger moral reason to blame non-hypocritically.

Now consider the data this provides. We have an example of a type of blaming – hypocritical blaming – where it is plausible, on one hand, that the blamer's standing is threatened, and where it is also plausible, on the other hand, that the blame, though *somewhat* valuable, is also *deficient* in value. Because of the factor of hypocrisy, the blame is not as valuable as it might be. Might this be more than a mere coincidence? Might one observation not explain the other?

One possibility is that Blake's hypocritical blame is deficient in value because he loses standing to blame Amy hypocritically. But that does not connect the two observations in an appropriate way. We seek to understand a particular kind of relationship between hypocrisy

and standing to blame: hypocrisy is one factor that seems to play a role in grounding a loss of standing. What we're after is an explanation as to why this is so – in virtue of what, exactly, does this factor play that role? Positing that hypocrisy loses value because hypocrites lose standing, then, only assumes the thing we are trying to explain.

Instead I will defend the converse proposal: Blake loses standing to blame hypocritically because hypocritical blame is less valuable. It is in virtue of this fact about hypocrisy – namely that it implies a loss in value – that it makes the person lose standing. I should be upfront about one limitation of this idea: it won't quite capture what people tend to say in response to the hypocrite and others who blame with deficient standing. I've never heard anyway literally say, 'Hey, your blaming me has less value!' But it will capture and justify moral concern about the blamer's loss in standing, and as we have already seen, the way we voice these concerns does not always represent the moral reality.

i. Homing in

However, we cannot ground standing simply in a loss in value. The loss in value that is relevant to standing is very specific: what is lost is the value of particular people engaging in accountability procedures with particular others. Let me explain.

Standing clearly cannot be affected by a loss in just any value. I might blame a colleague for insulting me when there is something else much better than I could be doing, such as treating a patient. But *that*, clearly, cannot be a reason why I would lose standing to hold my colleague accountable, for the loss here has nothing to do with blaming at all. The idea of standing we are looking to explain is, after all, standing *to blame* a person. If losses in standing are related to losses in value, it looks like these losses must be closely connected to the value of blaming in particular, rather than losses in any possible set of values.

But this is still not specific enough. We can blame others in a way that prevents them from responding appropriately to wrongdoing because of the circumstances: it might be that if I blame my colleague for insulting me, another colleague will intervene and help her laugh it off, whereas if I don't blame her, she's more likely to think about what she did and apologise.

But, again, that sort of consideration does not affect my standing. Recall how the ‘you can’t...’ responses I listed earlier are directed specifically at the person doing the blaming. That person is meant to lose standing because of some set of facts that are particular to *him*, and this set of facts is meant to trouble us about *his* blaming. This type of complaint does not arise when someone else will lead a person astray if I blame that person. There, the loss in successful correction is not importantly connected to the fact that it’s me who blames my colleague, and the manner in which I blame them. So, we should add this particularity condition: losses in standing are not quite grounded in the value of *blaming* in virtue of facts about the blamer. Rather, they are grounded in the value of *this particular blamer blaming* in virtue of about that blamer.

Even this, though, doesn’t fully capture the sort of concern that standing seems to raise. Some facts particular to me can undermine value without threatening standing at all: the fact that I have a certain accent, and my colleague is unresponsive to this accent and so would not take me seriously, is no threat to my standing, even though my accent is particular to me. So it matters not only whether the value of blaming depends on the particular person blaming, but also the reason for that dependence: the way in which the value, or the absence of it, is produced. The connection between the blamer and losses in value, in other words, cannot be purely instrumental – it cannot just be a causal connection between facts about the blamer and blamee and the success of their interaction.¹

The relevant loss, I will argue, is *procedural*. On this view, the *process* of a particular person holding another particular person accountable, for a certain kind of behaviour, can have a value or disvalue in and of itself, and can have a special kind effect on the value of the outcomes of blaming. Standing norms are about *that* type of value, and losses in standing track the extent to which blame fully instantiates it.

¹ Most of the above speaks against Gerald Dworkin’s view (2000), which explains standing by appeal to its instrumental effects. See Cohen (2013, Ch. 7) for other criticisms of this view.

ii. *Back to hypocrisy*

To see how this idea works in the case of hypocrisy, let us state it somewhat more formally:

Procedural Value: B (blamer) has defective standing to blame A (blamee) for X because, and insofar as, in virtue of facts about B, the procedure of B holding A accountable for X is deficient in value.

As indicated above, this assumes that the procedure of B holding A to account for X can be valuable in itself, because of some fact about B. What could make that true?

I think the answer depends on there being a morally important connection between B, on one hand, and other components of the accountability procedure, on the other: viz. A (the blamee) and X (the conduct that is the subject of the blame). Such *connective facts* about the blamer are the natural answer to why accountability between *those* people about *that* behaviour might have a moral valence in itself.

On this view, to understand losses in standing, there are two crucial things to consider. First, how do certain facts about the blamer connect her to: the person she's blaming, to the wrongdoing she is blaming the person for, or indeed some combination of these things (for instance, the fact that this person is guilty of the wrong)? Second, do these facts make a distinct difference to the value of blaming – a difference which is not entirely dependent on the outcomes of blame?

This rather abstract notion will gradually become clearer through looking at a range of concrete examples. But we can start with the by-now-well-trodden example of Supremacists. In Supremacists, the crucial connective fact seems to be the blamer's connection to the wrongdoing being addressed. What connects Blake to the racism for which he blames Amy is that he himself guilty of casteist abuse, which is a morally similar wrong. But that connective fact, I argued, implies a further connective fact: that Blake and Amy are in the same moral boat – they both confront a similar set of corrective demands. Through being connected to the wrong in a certain way, Blake is also connected to Amy in a certain way.

According to Procedural Value, the thing to ask is: In what way is this set of connective facts morally important? Why would Blake and Amy being in the same moral boat

make a difference to the value of *his* holding *her* accountable *for racism*? Of course, in this case, we already have an account. The fact that Blake is a similar wrongdoer makes possible a solidaristic connection with Amy which full collaboration about their similar wrongs creates, as well as any important moral insights about their wrongs which collaboration between them would help both parties learn. That package of things makes accountability between these particular people a good thing in itself.

Now, this doesn't yet explain why Blake's hypocrisy would have a *negative* effect on his standing: it only tells us what could be especially good or valuable about the blaming interaction. What we also need to know is why, in spite of these good connective facts about Blake, hypocrisy is a bad connective fact; one that implies a loss in value. But again, we already have an account in this case. Though he is guilty of similar wrongdoing, the special values that this fact could achieve are diminished by the fact that he fails to acknowledge his guilt. That is a bad relational fact about him, because it limits the values of the good relational facts about him: it inhibits full collaboration about wrongdoing, solidarity, and so on.²

In sum, whilst there are facts particular to similar wrongdoers that makes accountability between them procedurally valuable, hypocrisy attenuates that value. That, I'm arguing, is why hypocrisy diminishes standing.

² One implication of this thought may seem surprising. In standing terms, we might look upon a blamer who is guilty of similar wrongdoing more favourably than a blamer who is innocent of similar wrongdoing: if a similarly guilty blamer acknowledges that guilt, they can achieve certain procedural values that the innocent blamer cannot. But, this becomes less surprising once we see that the fact that the guilty blamer can achieve something procedurally better *also* means that their failure to achieve that better (by failing to acknowledge guilt) is more serious; and in this sense, the loss of standing of a hypocritical blame is more morally serious. The hypocrite's own loss is due to what he does, where the same is not true of the innocent blamer: the latter is not a similar wrongdoer, and so not responsible for failing to mutually deliberate with a similar wrongdoer.

2. Entrapment and Estrangement

One benefit of Procedural Value is that it illuminates a loss of standing in a broader range of cases where hypocrisy is not the core issue. In the next two chapters, I will explore more examples, but I will focus on two here by way of preliminary defence.

i. Reactive Accountability

Here is one instance of blame that looks deficient:

Entrapment. Blake is innocent of wrongdoing similar to racism. He is about to go to the pub with Ciara. Blake believes that if he invites Amy along, the issue of Ciara's race will come up, and Amy is liable racially to ridicule Ciara under those circumstances. Unbeknownst to Amy, he invites her in order that he can then blame her to her face for her racial abuse.

Given his plan to entrap Amy into wrongdoing, Blake plausibly loses standing to blame her for it. Why?

Again, Procedural Value asks: What is the relevant connective fact about Blake here, and why does it matter? Some might suspect that the relevant fact is the same as in Supremacists, because Blake, by inviting Amy to the pub, is partly responsible for the wrong she goes on to commit. In this view, although Blake may have been innocent of similar wrongdoing in the past, his entrapping Amy means that he is now guilty of similar wrongdoing; or, rather, he bears responsibility for the exact same wrongdoing. Thus, he blames Amy hypocritically.

However, hypocrisy is not the central explanation of this case, as the responsibility of the entrapper is not really the same as the responsibility of the person committing the wrong. To see this, suppose Blake takes the following as a reason to entrap Amy: he thinks his blaming her for the racist abuse will help to ensure that she avoids such abuses in future. He invites her along, though, because he wants it to be *him* in particular who gets to achieve this end, as he can then be a witness and be warranted in holding her accountable. In that case, Blake's connection to Amy's abuse seems very different from her own connection to her abuse – Amy acts wrongly because she fails to be moved by anti-racist values, whereas Blake sees those

values as a reason to hold her accountable. The idea that they are together in the need to correct for racist abuses is then not so compelling. However, it remains compelling that his standing to blame her is deficient, given that he deliberately lures her into committing the wrong *in order* to blame her for it. What seems distinctive of such cases, then, is that independent connection that Blake has to the wrongdoing.

Procedural Value can explain why that connection undermines standing, as follows. When blaming a person is the goal of a plan that requires getting them to act wrongly, the process of blaming them is non-ideal. In Chapter 2, I said that a core value in blaming a person is to try and help that person correct for wrongdoing by engaging them. That engagement, though, is especially important as a *response* to wrongdoing: wrongdoing is what creates the most stringent corrective demands, and thus what creates the most urgent need to reflect on how to comply with them.

Ideally, then, accountability processes are valuable as a *reactive* measure: it is good to help perpetrators (and victims) address wrongdoing *given that* wrongdoing has been committed. In contrast, blaming interactions are less valuable as a *proactive* measure. Indeed, the goodness of helping agents address wrongdoing is normally a *perverse* reason to have them commit wrongdoing in the first place.³ If a core value of blaming is to ensure that we are opposed to wrongdoing for the right reasons, using the means of wrongdoing to get there is anathema to that value.

So, when Blake blames Amy in Entrapment, he executes an intention to hold her accountable in a way that gets the ideal purpose of holding her accountable the wrong way around. He should blame Amy because she committed the abuse, rather than get her to commit the abuse because he can then (accurately) blame her. The fact that the procedure executes that perverse intention makes the blame less valuable in itself. That loss, in turn, explains the entrapper's loss of standing.

³ A compatible, but different view is that the problem with entrapment lies not in a perversion of accountability, but in the subversion the wrongdoer's moral capacities (Howard, 2017).

ii. *Value-constituting accountability*

A second instance of problematic blaming arises when the wrongness of the blamed conduct depends on our involvement in certain relationships with others. Consider:

Estrangement: Blake is a total stranger to romantic partners Amy and Ciara. Over several months, he happens to overhear discussions between the pair in a café, from which he learns that neither of them would be okay with the other being romantically involved with someone else. One day at the café, Blake witnesses Amy romantically engaged with someone else. He blames Amy to her face for cheating. Blake is innocent of wrongdoing similar to cheating, and is not disposed to commit such wrongs.

Blake's standing to blame Amy again seems deficient. But, again, the problem is not that he blames her hypocritically, as he has not done anything like that himself. The relevant connective fact instead seems to be that Amy's cheating is 'none of Blake's business'. The issue here is not simply about Blake's connection to the *wrongdoing*, then, but rather his estranged connection, or the absence of an intimate connection, with the *wrongdoer*, Amy (and indeed the victim, Ciara). Why might this different connective fact matter?

One answer draws on epistemic considerations. Given that Blake is not party to the relationship at all, we might think he is less likely to understand the nature of what Amy did, and what the consequences of blaming her might be.⁴ However, this cannot be the full story. It is typically true of strangers in Blake's situation that they lack knowledge of the situation, and the characters involved; indeed, to understand fully the character of strangers' relationships, we'd often have to do objectionable things, indeed things which violate their rights, such as stalking or spying. But this need not be the case. It is possible for strangers to

⁴ See McKiernan (2016). McKiernan adds that a norm of standing would be useful because blame by people who don't know the wrongdoer, and how they are likely to react, will generally lead to bad consequences (2016, p. 151). I bracket this part of her argument: as I've suggested, and will reiterate below, the possible bad consequences of blaming are not enough, by themselves, to explain losses in standing.

know the relevant facts, and where they do, this does not seem to have an impact on their standing. We can see this by imagining that Blake, though he is a stranger, is also an incredibly astute judge of others' behaviour, and thus knows even more than Ciara does about Amy's reasons for cheating on Ciara. This, intuitively, makes no difference: intuitively, this would not give Blake standing he otherwise lacks, and would not mean that Ciara now lacks standing to blame Amy.

The broader problem with grounding standing in knowledge is that knowledge is ultimately a matter of whether blame is *warranted* – a matter that, I argued in Chapter 2, is best kept separate from other ways of assessing blame. Standing looks like another type of evaluation that we ought to keep separate from warrant. As we saw with the examples of hypocrisy and entrapment, concerns about standing generally do not track whether or not the blamer accurately blames. Entrappers and hypocrites may know a great deal about the wrong in question, but their standing is nonetheless deficient.

Furthermore, differences in knowledge about the circumstances of wrongdoing do not explain what is unique about cases like *Estrangement*. In principle, we can lack knowledge about all types of wrongdoing others commit. We take wrongs to be committed all around the world without doing too much research into the nature of those wrongs, or into the consequences of blaming perpetrators. But the unease about holding perpetrators accountable with inadequate knowledge (e.g. for something they are rumoured to do on Twitter) seems very different from the unease about Blake's blame in *Estrangement*. So, to appropriately explain why Blake's lack of acquaintance with Amy (or Ciara) would threaten his standing, we should aim to explain why his lack of acquaintance matters in a way that does not reduce to an accuracy concern.

One popular suggestion along these lines is the following. Blake lacks standing as the stranger because his blame threatens Amy's and Ciara's ability to self-determine their relationship, and we ought to protect people against such threats by having certain

accountability norms in place.⁵ If our social norms permitted strangers to blame partners for wrongs like infidelity, partners would constantly have to worry about external influence, and this makes it harder for them to set standards for themselves about how they are to relate. In contrast, a norm that prohibits strangers holding one another accountable for infidelity makes it easier for partners to control this aspect of their lives. Since partners have a powerful interest in that control, we should prefer that the standing norm be upheld.

It may be plausible that relationship self-determination is on the whole better protected with a norm that reserves accountability for some and not for others. However, I don't think this is a fully satisfying account.⁶ It tells us why it would be *good* for everyone to endorse and promote a lack of standing in strangers; but that, it seems, is a different question from *why* people actually lack such standing in the first place. Some may think that the second question can only be answered by the first, but it doesn't seem that way to me.⁷ As indicated, we have all sorts of ways to evaluate the bad results of blaming. The point of standing assessments is much narrower. Standing is related to the goodness and badness of accountability processes in themselves, rather than the long-term effects of an accountability practice. It is of course important to protect people against threats to self-determination – we should avoid prying without consent, we should not spy on people, we should give people some leeway to form their own public self-image, and so on. But what Blake does is not just another threat along these lines. To see this, we can imagine that the relevant explanatory

⁵ See e.g. Siem (2019); Herstein (2020); Edwards (2018).

⁶ This argument may have other problems. For example, the mere fact that partners have a strong interest in self-determination does not establish that a norm of standing which prohibits accountability by strangers is justified overall. But, my point here is that this this is the wrong kind of argument to explain standing.

⁷ Compare Enoch (2009), who considers whether the type of argument – ‘wouldn't it be nice if p, therefore p’ – can work for certain axiological claims: e.g. ‘it is good that (it is good that I'm healthy), so it is good that I'm healthy’. That may be a valid inference. But, it is a different thing to show that this works for deontic claims: ‘it is good that (I have a reason not to blame), so I have a reason not to blame’.

factor is in place: there is a well-functioning norm against intrusion, so that people are scared to intrude, and know that if they try, they will be batted off and ignored. Thus, private control of relationships is securely protected. Here I take it that Blake *still* lacks standing to blame Amy as the stranger. This further illustrates that there's something about accountability for cheating between strangers that is in itself morally troubling.

Here, then, is my proposal. Strangers have defective standing to blame because special accountability between partners not only promotes value; it also constitutes value. It isn't just good for people to have special accountability for infidelity because this leads to more autonomous relationships. It is also *part* of the value of those relationships. Without accountability between partners for wrongs like cheating, the value of partners' relationships is in itself diminished. The reason Blake's standing is diminished is that his blaming Amy for infidelity cannot achieve that particular value. It does not constitute a valuable relationship between strangers.

This argument gains support from the idea that the value of certain relationships depends on special duties they involves. There are some duties we have (or owe with special stringency) because we are related to others in a certain way, such as professional duties between colleagues, and duties of care owed to family, friends, and other intimates. One basis for these duties is that they are constituents of the value of these relationships: unless friends make special efforts to care for one another, for instance, the friendship fails to realise its full value, as does a collegial relationship where colleagues fail to make sacrifices for one another's sake. The same story can be given for partners: unless partners make sacrifices as part of a mutual commitment – at least where it is mutually understood that both partners want this – the value of a partnership is threatened. This can ground a special duty between partners not to cheat on one another.

⁸ For a helpful discussion of this view, typically attributed to Scheffler (2001) and Raz (1989), see Renzo (2014).

Now note that for the same reason, special duties come hand-in-hand with special accountability: the value that grounds special duties *also* grounds special accountability for non-compliance with those duties.⁹ Friends and colleagues holding one another accountable for failing to comply with duties that a good friend or colleague owes, for example, is also part of what makes these relationships valuable. In other words, relationships where people do not account to one another for violations of their value-constituting duties are in themselves less valuable than relationships that involve the relevant accountability practices. And we can say something similar about partners: just as avoiding infidelity can be part of a valuable partnership, partners holding one another accountable for wrongful infidelity also constitutes a valuable partnership.¹⁰

Once we appreciate this value-constituting role of special accountability, we can see how Procedural Value offers a compelling explanation of Estrangement. A stranger like Blake cannot instantiate accountability that constitutes the value of special duties between partners, as only those who bear the special duties (– i.e. only Ciara and Amy –) can instantiate that value. This doesn't deny that Blake can *promote* the value by holding Amy accountable, for instance where his doing so encourages Amy to open herself to accountability with Ciara. But

⁹ Elinor Mason (2019, pp. 191-196), David Owens (2012, pp. 38-40), and Seth Shabo (2012) also defend the idea that blaming constitutes valuable relationships, though they spell this out in a different way, arguing that valuable relationships depend on blaming emotions like anger, resentment and remorse: without these interpersonal emotions, relationships we deeply value would be unrecognisable. (For a similar thought, see McKenna, 2012b, pp. 135-138). I am not sold on this. The fact that a relationship wouldn't be as it is without certain emotions does not show that these emotions are a necessary part of what is valuable about that relationship. I tend to think that accountability practices are the core value-constituting feature. The negative blaming emotions which often accompany accountability practices are like impurities in the boiling water: although it is barely possible to cook a tasty pasta without them, they are not part of the value of the meal. Compare Pereboom (2001, ch. 7).

¹⁰ Echoing the discussion in the last subsection, I think this is so on the condition that infidelity has happened – it is not a reason to practice infidelity in the first place!

the process of *him* holding Amy accountable cannot in itself instantiate this value; only the process of accountability between partners can do this. This deficiency in the value of the procedure that Blake initiates, then, explains why his standing is defective, in a sense that Ciara's is not.¹¹

To recapitulate: I have argued in this section that hypocrites lose standing because they blame in a way that is procedurally deficient, where that deficiency is due to the hypocrite's failure to acknowledge similar wrongdoing. In support of this, I have pointed out how we observe similar procedural deficiencies in other standing cases, like Entrapment and Estrangement. However, this is not the only possible explanation for why hypocrites suffer a loss in standing. In what follows, I will offer a comparative defence of Procedural Value, showing that it overcomes the disadvantages in various alternative views, whilst also being able to explain their attractive features.

3. Ledger and Commitment

Let us start by exploring two alternatives: first, that hypocrites lose standing because of their bad behaviour, and second, that they lose standing because of their bad character.

i. Ledger

Perhaps the most instinctive view about hypocrites' standing is as follows: standing is explained not by the value or disvalue of blaming, but by the value or disvalue of the blamer's

¹¹ It might be that relational self-determination still has an important role to play here. We might think, for example, that accountability between partners is part of relational self-determination: accountability is one way that partners get to set the terms of their interaction, so special accountability between partners constitutes the value of a self-determining partnership. On this view, Blake loses standing because unlike Ciara, his blaming Amy does not instantiate the value of relational self-determination. This may be an appealing thought; but if it is, its appeal supports Procedural Value, as the focus here is on the value of the procedure of accountability *in itself*, rather than the value that a practice of accountability promotes in the long run.

past. In Supremacists, Blake has a bad past – he is blameworthy for casteism – and that bad past is what threatens his standing to blame. We can call this view *Ledger*. Roughly, it is the idea that people with good moral track records get to hold others to account, and people with bad moral track records do not. However, this cannot be right. Although the badness of a person’s past can matter, it does not determine standing by itself.

Why endorse Ledger? It might be tempting to support it by appeal to a principle of desert. Some believers in desert might think that that we ought to ensure that benefits and burdens are distributed in a way that prioritises the interests of non-blameworthy people over the interests of blameworthy people.¹² We might think this supports a principle that favours blaming by the non-blameworthy rather than by the blameworthy. A possible view is as follows. As blaming is beneficial for a blamer, and burdensome for a blamee, blaming by the non-blameworthy of a blameworthy person is good, as this benefits someone who is not blameworthy whilst also burdening someone who is blameworthy. In contrast, matters are not quite as good, from a desert perspective, when a blameworthy person blames another blameworthy person: that does burden one blameworthy person (blamee), but it also benefits another blameworthy person (blamer). That is one way to argue for a preference against people with bad ledgers doing the blaming.

I have already suggested that blaming can be a big benefit to the blamee, so it might not be the best way to deliver any requirements of desert, if any. But equally, it is not clear that blaming is always a net benefit to the *blamer*. There are significant benefits, such as in moral expression and normative engagement with others; but we’ve also seen that there can be costs big enough to release us from general duties to blame, such as reaction costs, and the fact that blame can be draining, and can put a strain on relationships.¹³ More importantly, even if we

¹² This thought features more prominently in desert views of punishment: Compare Wellman (2017), Husak (2011); Kessler and Ferzan (2009); Moore (1997). For critique, see Tomlin (2014), Tadros (2011) and Dolinko (1991). (Also compare Chapters 2 and 4 for examples of desert views about blame.)

¹³ Compare Shoemaker & Vargas (2019).

allow that blaming is typically a benefit to the blamer, this is not relevant in assessing the blamer's standing. For example, Blake may blame Amy hypocritically, but in a way that ends up being extremely costly for him overall. It seems clear that his hypocrisy still diminishes his standing to blame her, though any desert consideration against him blaming cannot tell us why. If anything, desert considerations yield the opposite result: Blake is blameworthy and blame is burdensome for him, so we should prefer that he blames rather than that a non-blameworthy person blames.

Procedural Value does not face this problem. The value of a blaming procedure is the thing that determines standing, regardless of whether this value is overall valuable for the blamer, or the blamee, is whether it achieves a good distribution of benefits and burdens. As the previous two chapters have outlined, benefits and burdens of blaming do matter in other ways, such as to whether certain acts of blame are permissible, but they do not directly affect a blamer's standing.¹⁴

A second disadvantage Ledger is that it struggles to explain why *similarity* in wrongdoing is relevant to standing. Our moral ledgers can be bad, yet that badness can have little to do with the type of behaviour that we blame others for. For instance, imagine that in Supremacists, Blake blames Amy for racism, having himself cheated on his partner, or committed several acts of petty theft. Whilst these things have an impact on Blake's ledger, they do not intuitively diminish Blake's standing to blame Amy *for her racism*, given that the wrong of racial abuse is different in nature. At least our assessment of Blake's standing to blame Amy for racism when he commits casteism is different than it is when he commits a non-similar wrong.

Procedural Value naturally tells us in what way similar wrongdoing matters.¹⁵ Similar wrongdoing is one fact that enhances the value that a process of accountability can achieve:

¹⁴ The next chapter says more about the relationship between standing and permissibility.

¹⁵ How do we assess standing to blame between wrongdoers whose wrongs are not similar either in gravity nor in nature? Example: why do I seem to have defective standing to blame someone for petty

given similar wrongdoing, accountability processes can realise the value of solidarity, where this involves mutually recognising that parties confront urgent moral demands that are similar in nature, and exchanging thoughts about this common problem. As conceded in Chapter 3, bonds of solidarity are available between non-similar wrongdoers, but this particular type of bond is not at stake when Blake blames Amy for racism as the cheater or the petty thief.

ii. *Commitment*

A different account of standing, proposed by Patrick Todd (2017), is better-placed to respond to this problem. Todd's view – let's call it *Commitment* – is that hypocrites lose standing because (and only because) they are not committed to the values that they invoke when blaming: i.e. the values flouted by the blamed actor.¹⁶

Commitment does not imply that just any bad past impairs standing; indeed no bad past is sufficient to do this on its own. Rather, the relevance of Blake's particular bad past – his guilt of casteism – is that it provides a signal of his moral character: it signals that he does not genuinely care about the anti-supremacist values Amy flouted. Rather than his bad

theft, if I have committed mass atrocities? The unhelpful answer is, I am not entirely sure. Some may see this as a case of hypocritical blaming. My sense is that standing and hypocrisy in a blamer who is more or less 'beyond the pale' is not explained in the same way as it is cases like Supremacists. Normally it will be hard for someone with a past so bad to be willing to reflect on her wrongs, and to be sensitive to the value of accountability: this may be a source of procedural disvalue. But, where such a person *does* manage to reform, perhaps there can be procedural value in her acknowledging her wrongdoing: doing this can help the blamee recognise that correction is an enduring possibility, even when a person seems beyond the pale. Perhaps there is also procedural value in the blamer expressing this fact to others – for example if a mass-atrocity-committer has stringent duties to prevent others from spiralling into serious wrongdoing.

¹⁶ Also compare Rossi (2018). Commitment will also struggle to explain gravity (see previous note): an atrocity-committer may have defective standing to blame a petty thief, even though the atrocity-committer would never steal from a shop.

behaviour, it is that more basic fact about his character which takes his standing away. If that signal turns out to be misleading – if for example Blake has radically reformed his attitude toward anti-supremacist values – Blake would suffer no loss in standing to hold Amy accountable.

So, similar wrongdoing is relevant, according to Commitment, because the moral values flouted by similar wrongdoing overlap with the values that hypocrites invoke when blaming others. This, Todd argues, is what differentiates Supremacists from the case where Blake cheats on his partner, or is a serial petty thief: the values flouted by those acts do not overlap particularly with the values flouted by racism. Those past mistakes are then not good evidence that Blake lacks commitment to the values that he blames Amy for flouting.

This is an improvement on Ledger, and I suspect commitment to moral values bears relevance to a person's standing. But as I'll now argue, both Ledger and Commitment fail to deal appropriately with the other cases that I have discussed.

iii. Back to the other examples

We might wonder whether the other cases – Entrapment and Estrangement – require a similar explanation to the case of hypocrisy. Before I consider that worry, let us explore why they pose a problem for both Ledger and Commitment.

a. Entrapment revisited

As we observed when discussing Entrapment, an entrapper might lure a wrongdoer into acting wrongly because he takes certain values as a positive reason to hold this wrongdoer accountable. And, where the plan will be effective in protecting those values, such a plan might be justified overall. Thus, an entrapment plan need not taint a person's moral ledger. So Ledger does not deal well with this example.

Entrapment is also a problem for Commitment. If Blake takes the need for Amy to reflect on her anti-racist values as a reason to blame her, that is powerful evidence that Blake is committed to those values. But intuitively, his standing is defective nonetheless. It might

seem that the entrapper's intention to be the one who gets to hold the blamee accountable is countervailing evidence: such an intention may indicate that Blake cares more about his getting to blame than he cares about addressing prejudicial abuse. But this needn't be true. It may be that Blake first and foremost cares about the values. For example, he might be certain Amy will commit racial abuse, but uncertain that there are going to be other circumstances in which someone will be warranted in holding her accountable for doing so. Indeed, Blake might be certain that his entrapment plan will *reduce* the amount of racial abuse Amy is going to commit overall. All of this would suggest that Blake is committed to the relevant values. However, because he forms that plan, Blake's standing is still deficient, given that the plan involves the further goal of himself being the person who gets to hold Amy accountable.

b. Estrangement revisited

Similarly, both Ledger and Commitment will struggle to explain the loss of standing in Estrangement, when Blake, as the stranger, blames Amy for cheating on Ciara.

In Estrangement, there is no prior act of wrongdoing, so the contribution of some prior bad act to Blake's ledger cannot be a basis for his loss of standing. One response is that Blake taints his ledger *in the act of blaming*, as blaming as a stranger is itself morally bad. However, as we saw, blaming, even as a stranger, might be a good thing for Blake to do overall; it may end up being effective in getting Amy to reform, and to make things up to Ciara if she is willing to entertain that possibility. Even so, Blake's standing is deficient in this case.

Neither does Commitment deal well with this example. Recall that Estrangement assumes that Blake would never himself cheat. Indeed the reason for his preparedness to intrude in Amy's affairs might be precisely this: he cares so much about the value of fidelity that he's willing to intervene in others' affairs to advance it. But he lacks standing to hold Amy accountable all the same.

c. Basic Moral Standing?

Some have argued that the cases of hypocrisy and estrangement are very different, and thus do not require a similar explanation. So perhaps what I have just said is not a problem for alternative views. I am not persuaded by this response, however.

Consider Todd's argument (2017) for treating hypocrisy and estrangement as different cases of standing. Todd notes that, as the stranger in the café, Blake does not lose standing to simply *form an attitude* of blame towards Amy for cheating on Ciara – if he formed such an attitude about Amy's cheating, we wouldn't object on the grounds that he is a stranger. We would object, though, to Blake *expressing* this attitude to Amy about her cheating. However, matters are different for Blake as the hypocrite. Unlike the stranger, we would object to the hypocrite simply forming an attitude of blame toward Amy for her racism if he is himself prepared to flout anti-supremacist values. On this basis, Todd says:

Because one lacks standing (with a person) to *express* blame, it does not follow that one lacks standing (with morality) to *feel* blame [...] I will call the standing to *feel* blame the 'basic moral standing' to blame [...] (2017, p. 350)

This might seem to suggest that the type of standing that hypocrites lose – i.e. 'basic standing', or standing 'with morality' – is fundamentally different from the type that strangers lose – i.e. 'non-basic standing', or standing 'with a person'. If that is right, then the reason why they each would lose standing may be fundamentally different.

But this is not a powerful response. For one thing, I doubt this contrast generally works: I doubt hypocrites, or indeed hypocrites who lack commitment to certain values, have no standing to feel blame about others flouting those values. Imagine, for example, that Blake hits Ciara for fun, revealing his lack of commitment to values against harming people for fun; and next week, Amy hits Blake for fun. It is not intuitive that Blake has no standing to feel blame towards Amy for hitting him – to be upset at her, for example. I say more in the next chapter about why this is so, but note for now that it seems disturbing to think that victims of wrongdoing do not have basic moral standing to feel wronged, even if they have committed similar wrongs and are disposed to committing similar wrongs. Morality cannot be so heartless as to condemn victims' warranted feelings on those grounds.

For another, the distinction between standing to feel and express blame does not show that the two issues of standing have nothing in common; nor does it explain why one (standing to feel) is basic, whereas the other (standing to express) is not. Indeed, there are equally strong reasons to think that expression is the basic thing. One is that in both cases – where Blake is the hypocrite, and where he is the stranger – expressing blaming is objectionable. It's true that only in the hypocritical case do people also feel uneasy about merely attitudinal blaming: as Todd points out, some might even feel uneasy about people hypocritically blaming fictional people, such as a serial killer forming blaming attitudes towards King Joffrey, a *Game of Thrones* villain, for something he does in the fantasy series (Todd, 2017, p. 364). It is a merit of Commitment that it captures this unease. Yet, since expressing blame to others is the thing that seems objectionable in *both* cases, why not think expression is the core moral concern? Part of the argument seems to be that standing 'with morality' is basic, and the hypocrite lacks standing 'with morality' to blame; whereas standing 'with another person' is not basic, and the stranger only lacks standing 'with another person'. However, standing 'with morality' and 'with other people' surely overlaps. Morality, after all, is largely about how we treat other people; so this offers no reason to deny that the standing to express blame is central.

The centrality of blame that is expressed through an accountability procedure is also supported by two observations. First, we saw that standing norms have an *interpersonal* flavour: it is about particular people blaming particular others for particular things, and how all of these components of a blaming interaction connect to each other. The communicative paradigm – of engaging others about wrongdoing – more fully engages that interpersonal structure of the norm. In the communicative paradigm, we can consider how a person is treating another, and what type of relationship is being built or thwarted through that treatment. Such issues are not engaged when a callous person blames King Joffrey.

Second, it is extremely difficult to avoid feeling blame.¹⁷ In the previous chapter, I suggested it will be difficult for a person to engender commitment to a moral value, and, unless they have that commitment, to ensure they don't communicatively blame others for flouting that value. That difficulty is magnified when we extend this norm to govern attitudinal blaming. It is even harder to avoid so much as *having* negative attitudes towards someone who wrongs us, or wrongs someone we care about, even if our own character is such that we are disposed to committing similar wrongs. In contrast, expression is more clearly under our control: there is a step between having a blaming attitude and making a decision about whether, how, and to whom to express that attitude. Since communication is more under our control, and since moral norms should be somewhat sensitive to what we can control, communication is a more appropriate candidate for a norm of standing.

Procedural Value, of course, explains the centrality of communication – i.e. why the hypocrite lacks standing *both* 'with morality' *and* 'with people' to communicatively blame them for certain wrongs. According to Procedural Value, standing matters, not only as an internal issue in the blamer (for instance, an issue about his poor character), but also from the point of view of the person blamed, where that person has interests in being blamed in procedurally better rather than worse ways. This would justify concern about certain cases of attitudinal blaming as well. Certain blaming attitudes suggest that a blamer is disposed to blame actual people badly, and will communicate the blame badly to those people. My thought is that this a much more central moral concern than any intrinsic disquiet about my blaming King Joffrey.

4. Detachment

Another disadvantage of both Ledger and Commitment is that both views confront the sceptical challenge to the significance of standing I described at the beginning of this chapter. The challenge is to explain how a fault in the *blamer* affects our moral assessment of the act of *blame*.

¹⁷ Thanks to Johannes Roessler for discussion of this point.

As Macalester Bell (2012) puts the worry:

Hypocrisy is a moral fault [...] In fact, people may, and frequently do, evince a wide variety of moral faults through their blame: they can show meanness, pettiness, stinginess, arrogance, and so on. But while people may manifest hypocrisy and other faults in their critical interventions, there is no reason to conclude that these faults always undermine a person's standing to blame (2012, p. 275).

Here is how I interpret what Bell is getting at.¹⁸ The right way to understand blame which evinces faults – such as hypocrisy, intentions to entrap, intrusions, and so on – is to ‘detach’ our assessment of those faults from the act of blaming itself. We may have independent moral reasons not to be hypocrites; not to intrude in others' lives; not to form problematic intentions, and so on. However, why should these reasons bear on the question of whether and how we should blame others – why think that they are appropriately ‘attached’ to our evaluation of blaming?

This worry applies forcefully to Ledger, which focusses on the blamer's past bad behaviour. It is true that the blamer's past was morally bad. But, the sceptic will ask: by what moral sorcery does that bad past, in and of itself, affect the morality of her present act of blame? It seems more appropriate to assess the two things separately – the hypocrite's past might be morally bad, but her blaming itself might be morally good.

The detachment worry also applies forcefully to Commitment, which, rather than focussing on the blamer's bad past, instead focusses on her bad character (in particular, her non-commitment to certain moral values). Again, the sceptic will ask: By what moral sorcery does that defective character, in and of itself, make the act of blaming morally bad? Indeed, blaming may turn out to address the blamer's own non-commitment, and might help rectify the fault that is meant to count against blaming. So again, it seems more appropriate to assess the two things separately – the blamer's character might be bad, but the blaming itself might be good.¹⁹

¹⁸ This is similar to Lang's (forthcoming) interpretation.

¹⁹ Todd (2017) hints at the following response: ‘any criticism [a non-committed blamer] might direct toward [the blamee] would have to be [...] in bad moral faith’ (2017, p. 356). But, as we'll see in the

In contrast, Procedural Value does not face this challenge at all. The reason why certain facts about a blamer threaten standing is precisely *that* they hamper the moral value of particular acts of blaming. Of course, it needs to be shown that this is possible; else the sceptic will agree with Procedural Value, but deny that anyone actually has defective standing to blame. But I have already given a number of arguments for why certain facts, like hypocrisy, estrangement, and entrapment, can affect procedural value.

5. Acknowledgment

In both attitudinal and communicative cases of blaming, another core standing-relevant fact is acknowledgment. Where a blamer fully recognises and is willing (whether privately or publicly) to reflect on her similar flaws, it is intuitive that the blamer considerably restores standing to blame others for those flaws.²⁰ Again, unlike Procedural Value, Ledger and Commitment fail to explain why.

Ledger, for example, is about our bad past, and since acknowledging a bad past does not make it good; this fact should have no impact on standing. Perhaps a proponent of Ledger will argue that whilst our morally bad acts threaten standing, our morally good acts might restore standing. Thus, whilst it is bad for Blake to commit casteism, it is good for him to acknowledge casteism, and that good act restores the standing threatened by his bad act. However, as good as I've argued it is to acknowledge wrongdoing, it is, surely, much worse to act wrongly in the first place: it is much worse to commit a casteist wrong, for instance, than it is good to acknowledge committing a casteist wrong. So even on this way of understanding Ledger, acknowledgment cannot do much to restore Blake's standing.

next section, it is a mistake to think that such a blamer must blame in bad faith. For he can openly and fully acknowledge his lack of commitment to the values invoked in the act of blame. When he does so, the blame is not in bad faith..

²⁰ This is shared by proponents of the Equality account – see next footnote – as well as Herstein (2020) Edwards (2018); Roadevin (2018) and Rossi (2018).

It might seem that Commitment more easily captures the impact of acknowledgment. Part of Blake being committed to anti-supremacist values, some might think, is to acknowledge having flouted them. One argument for this is the following. Being committed to values we have flouted involves complying with our duties of correction for flouting them, and these duties involve acknowledging what we did. This may be right, but it does not fully explain the significance of acknowledgment. For it is also true that we can acknowledge our wrongs *without* being committed in our opposition to the values we flouted, and acknowledgment clearly matters in such contexts. To see this, assume that Blake is not committed to anti-supremacist values: suppose, for example, that he is disposed wrongly to make an offensive remark in future. Now consider two ways of publicly or privately blaming Amy. One involves not acknowledging his faulty disposition at all. The other involves fully acknowledging it, by saying or thinking, 'Look, I definitely cannot claim that I myself am not liable to do this type of thing again; I'm ready to think about or discuss this flaw in myself.' Intuitively, Blake considerably restores standing if he blames Amy in the second way, despite lacking commitment to anti-racist values either way .

Some might respond that the fact that Blake acknowledges his wrongdoing shows he has a more committed opposition, overall, to the values he flouted, and that is why his standing is restored when he acknowledges it. But even this doesn't capture the difference. Suppose, for instance, that the version of Blake who refuses to acknowledge his mistake is much less disposed to flout the relevant values himself, thereby making him much more strongly committed to those values, overall, than the version of Blake who acknowledges his character flaw. It remains intuitive that the uncommitted Blake, who acknowledges wrongdoing, restores his standing in a way that the more strongly committed Blake, who refuses to acknowledge, does not. Unlike Commitment, Procedural Value explains why. Though Blake is uncommitted either way – though his character may not be good – it is still *procedurally* better, for both this person and the others they hold accountable, to undergo a process of mutually discussing their similar flaws.

To be clear, the above is not meant as a rejection of versions of Commitment; it only rejects the idea that commitment is the core explanation of a hypocrite's loss of standing. Indeed, I think a person's commitment might be relevant to determining their standing, as it can have an impact on the procedural value of accountability. Commitment to moral values may motivate the blamer to acknowledge wrongdoing, for example. Similarly, witnessing the blamer's commitment may inspire the blamee to participate in mutual deliberation about wrongdoing. And engaging in accountability with a morally committed person, who is sensitive to the relevant values, may help ensure that the right set of substantive issues are discussed. Otherwise put, Procedural Value can help explain when and why commitment matters, in the following way. Having a *moral* commitment – a commitment to substantive moral values – often motivates a person to have a *mutually deliberative* commitment – a commitment to reasoning with others about responding appropriately to those values. However, these two commitments can also come apart. Where they do, the deliberative commitment has a clear procedural relevance, and thus a clearer relevance to standing.

Time to take stock. In these last two sections, I have illustrated the ability of Procedural Value, relative to other Ledger and Commitment, both to respond to the worry that facts about the blamer should be detached from the morality of blaming, and to explain why acknowledgment restores standing. However, a third alternative promises to meet both desiderata. Let us finally compare Procedural Value with this third alternative.

6. Equality

Several philosophers have suggested that hypocrites, in failing to acknowledge wrongdoing, lose standing because they deny the equality of the blamee.²¹ Let's call this view *Equality*. Equality fares better than Ledger and Commitment on the detachment score. But it fares worse

²¹ See Fritz and Miller (2018); Wallace (2010); Duff (2010, p. 127, pp. 138-139); Tadros (2009, pp. 394-404). For other criticisms of an egalitarian view of standing, see Lippert-Rasmussen (forthcoming).

than Procedural Value on the acknowledgment score, as the significance of acknowledgment is not simply that it secures a standard of equal treatment or regard.

Equality is roughly the argument discussed in Chapters 1 and 3: that the blamer who doesn't acknowledge his similar faults whilst blaming a person engages in an asymmetrical treatment of himself vis-à-vis the blamee, in a way that denies the blamee's equal moral status. But the conclusion here is not merely that the blamer's denial of equality makes the blame wrong; Equality also holds that this fact makes the blamer lose standing²².

Given that Equality already implies that there is something morally wrong in the blame, it seems clear that this account does not face the detachment worry. Equality can straightforwardly appeal to a relevant fact about the blamer – i.e. his denial of the blamee as his equal – to explain why his blame might be objectionable. I have of course suggested throughout that hypocritical blaming is not best understood as involving a denial of equality. But the point here is that if that *is* what's going on – if a hypocrite does indeed show a lack of equal regard for others – then it is plausible enough, where that fault is committed in a blaming interaction, that this affects the morality of the interaction.²³

Equality might also seem well-placed to account for the significance of acknowledgment. However, I do not believe that the central reason why acknowledgment restores standing, and why its absence diminishes standing, has to do with equality in the distribution of blame.

There are numerous possible accounts of the type of unequal treatment committed by the blamer who does not acknowledge wrongdoing. Recall, for example, Wallace's account: that by failing to acknowledge his wrongdoing, the blamer treats his own interests in avoiding

²² Again, for further explanation of this distinction, see Section Chapter 5, 1.i.

²³ Bell (2012) seems to reject this idea on the ground that unequal blaming might still have good consequences (2013, pp. 274-276); but again, that is beside the point. The core challenge in the detachment worry is to show why blaming itself can be objectionable in light of facts about the blamer, and Equality seems very well-placed deliver that result.

blame as weightier than the blamee's interests in avoiding blame. Does this fact explain the relationship between acknowledgment and standing? Well, suppose I blame someone for cheating on their partner, and when I do this, I remember that I also steal things from the shop, though I do not blame myself for stealing from the shop. On Wallace's argument, we should think that I lack standing to blame this person in just the same way as the ordinary hypocrite: after all, I treat my own interests in avoiding blame as weightier than the blamee's. But, intuitively, the source of the ordinary hypocrite's loss of standing is different. If I acknowledge my stealing from the shop, that would not have the same impact as Blake acknowledging his casteism when blaming Amy for racism.

Similarly, consider Fritz's and Miller's version of Equality (2018), which says that the hypocrite loses standing because he has a 'differential blaming disposition' in response to his own and others' violations of a moral norm, on the ground that this disposition 'rejects the impartiality of morality with respect to violations of [that particular norm]' (2018, p. 125). Why think that this disposition is especially problematic when it comes to violation of '*that* particular norm'? The basic worry is about partiality in how we treat others. But partiality in our responses to norm-violations is no worse when those violations are very similar in nature, compared to when they are very different in nature. Again, we can see this in the cheating versus theft example. If Blake is disposed to blame Amy for cheating without blaming himself for theft, he blames her partially. Again, this does not explain why this case intuitively differs from cases like Supremacists.

We can further exemplify this general limitation by thinking about differences, not only in whether people self-blame whilst blaming others, but also in the *manner* in which they self-blame. Imagine, for instance, that Blake is disposed to blame Amy in a proportionate though heated tone for cheating on Ciara; but when it comes to acknowledging that he also cheated on his partner, Djamel, starts to cool down a bit, and blames himself for cheating on Djamel in a somewhat calmer fashion. This is both a differential blaming disposition, and involves shielding himself from the type of blame to which he subjects Amy. But again, intuitively, these facts do not stop Blake's acknowledgment of wrongdoing from significantly

restoring his standing. Equality in the distribution of blame, in other words, should require equality not only in who is blamed, but in the tone in which they are. But that is too pedantic a requirement for having good standing to blame.

Once again, this is not to argue for the irrelevance of equality considerations altogether. Indeed, as with Commitment, Procedural Value helps delineate the truth in Equality. Equal regard between blamer and blamee is part of the terms on which a valuable discursive process between them may proceed. If blamer and blamee treat one another as different in moral status, each might be less inclined to listen and respond to the other, and to offer their own thoughts about similar wrongdoing. That has likely implications for the moral quality of the accountability procedure between them, and thus likely implications for the blamer's standing.

At the same time, though, equal regard does not fully capture how acknowledging similar wrongdoing affects standing, for a reason sketched in Chapter 3. Beyond being a way of showing that we are *not partial* to ourselves in our blaming responses, acknowledgment is also a way of showing that we are *together* with those we blame in responding to certain types of wrongdoing. A key part of acknowledgment's procedural value is its ability to create a solidaristic bond – one that relies not only on seeing one another as equally morally worthy, but also on seeing one another's moral situations as appropriately included in a mutual conversation about wrongdoing.²⁴

Conclusion

²⁴ I used to think a similar distinction applies to practices in parts of rural India, where certain communities look up to foreigners, seeing them as greatly superior, but nonetheless will not eat at the same table, or use the same kitchenware (though all the kitchenware would be of the same quality!). One understanding of the attitudes underlying such a practice is that whilst foreigners are not inferior, they cannot appropriately partake in this activity together with natives. Though I doubt this is empirically accurate, it suggests that some distinction along these lines is possible and important.

This concludes my initial argument for Procedural Value. The remainder of this thesis further defends its usefulness as an account of standing. Its usefulness, we have seen here, lies not only in capturing a broad range of standing-relevant facts, but also in meeting important limitations of popular alternative views, and, in supplementing those alternatives, by helping to explain the importance of the factors they identify.

But there's a notable omission in the discussion so far. Though I have claimed that standing makes a moral difference – it makes a difference to the moral quality of accountability procedures – I've said little about what difference this should make to our deliberations about what to do in practice. What *is* standing as a moral property, and how should considerations of standing guide us in determining how blame ought to be practised? The next chapter takes up these questions.

Chapter 6. Standing, on a Scale.

The previous chapter argued that standing to blame depends on particular facts about the blamer: facts that alter the procedural value of accountability. But what *is* standing – what *is* the thing that depends on those value-affecting facts? Without knowing what the standing to blame is, we cannot know what losses in standing mean for how we ought to blame others. We tend to think that if a blamer’s standing is defective, it follows that he acts wrongly in blaming, or that he lacks a moral authority to blame, or an entitlement to blame. Perhaps surprisingly, once we focus carefully on the facts that affect standing, it turns out that these views all rely on a mistake. Indeed, the standard way of thinking about the normativity of standing does not marry with a plausible view about how standing-relevant facts actually work.

To get a sense of the standard approach, let us review our indicative list of responses to defective cases of blaming:

‘You can’t talk to me about that!’

‘You don’t have a leg to stand on!’

‘You don’t get to blame me!’

‘You have no right to complain!’

‘You’re in no position...!’

These phrases imply that the person has *no standing whatsoever* to hold another person accountable. That idea is implicit in the categories each phrase refers to. Either we ‘can talk’ about something, or we cannot; we either ‘have a leg to stand on’ or we do not; we either ‘get to blame’; ‘have a right’ to blame; ‘are in a position’ to blame or not (and so on). Standing, then, is a bit like an on-off switch: a person either has full standing or none at all. So, a person who loses standing has none at all.

This approach – I will call it *Binary* – should be rejected. Standing is something we can have more or less of: a person who has standing may not have full standing, and a person who loses standing may retain some standing. I will call this alternative view *Scalar*.

Scalar is naturally supported by the Procedural Value account of standing-relevant facts: the procedural value of someone holding another accountable need not be an all or nothing matter, as some values can be scalar – there can be more or less value in something, so accountability might be more or less procedurally valuable. Thus, degrees of Procedural Value might explain degrees of standing. However, I will argue that Scalar is independently plausible. Whether we agree with Procedural Value or not, Scalar explains a greater range of central cases than Binary does, and explains them in a more compelling way.

Probably guided by the nature of our ordinary discourse on these matters (illustrated by the ‘you responses’ above), philosophers, too, are pre-reflectively inclined towards *Binary*. Sometimes this is because discussions do not explore in detail the normative implications of standing, instead focusing on whether particular ‘conditions’ are necessary and sufficient for a person to have standing.¹ This focus on conditions of standing naturally lends itself toward binary thinking: a tempting conclusion is that where certain conditions are absent, a person has no standing; and where those conditions are present, a person has standing. At other times a particular binary view (e.g. a permission, a power, or an entitlement) is assumed, because the thing that commentators aim to explain and defend is precisely our ordinary discourse on standing: as the ordinary discourse is binary, a binary view is naturally what people want to vindicate. Whether implicit or explicit, there is a marked trend of describing a blamer ‘forfeiting the standing’ to blame.² Again, this phrasing paints a binary picture: the blamer forfeits this

¹ Compare Seim (2019); Fritz & Miller (2018); Todd (2017); McKiernan (2016); Friedman (2013); Bell (2012); and Cohen (2013).

² Compare Radzik (2011, p. 590); Todd (2017, p. 371); Fritz and Miller (2018).

singular thing, ‘the standing to blame’, meaning she no longer has that thing. Similarly, the term ‘standingless blame’ is often used as shorthand for defective cases.³

That said, I am aware of no explicit argument for why we should be thinking about standing in a binary way at all. Of course, even in the absence of such an argument, the right view may still *be* a binary one. However, in Section 1, I give reasons to doubt this, by rejecting the three of the most familiar versions of Binary: first, that standing is a certain kind of moral permission; second, that it is a moral entitlement to blame; and third, that it is a moral power to demand a certain response from the blamee. Though perhaps credible at first glance, none of these satisfactorily explains core cases. In core cases, considerations of standing generally matter regardless of whether the blamer has or lacks a permission, authority or entitlement to blame.

A more promising view is that standing affects the *moral appropriateness* of blaming a person. In Section 2, I explain this idea. I also explain how an appropriateness view of standing supports Scalar – roughly, as blaming can be more or less apt, it should be no surprise that we can have more or less standing to engage in it.

Section 3 illustrates the intuitive appeal of a scalar view which holds that standing is an appropriateness standard. The first intuitive thought is that there are multiple standing-relevant facts, and many of these contribute, positively or negatively, to blame’s appropriateness. Through looking at the difference these facts make, we can see that, often, their contribution is not to ‘switch standing off or on’, but more plausibly to diminish standing or enhance it. A second intuition is that *each* standing-relevant fact is itself scalar in nature. Each such fact is something that can be more or less significant in a given case, and it is plausible that standing increases or decreases in correspondence.

³ Compare PEA Soup (2019); Tognazzini (2019); King (2019); Seim (2019); Fritz & Miller (2019); Todd (2017); Bell (2012)

Lastly, I consider a strategy for rescuing Binary whilst accepting both these arguments. We might accept that whilst standing-relevant *facts* work in a scalar way, *standing* is itself a binary category. I call this view *Threshold*, and give a general objection to it.

1. Standing: Permission, Authority, Entitlement?

Before raising problems with binary views, let me explain in a bit more detail the basic idea that I will reject.

In the previous chapter, I offered three examples of standing-relevant facts: hypocrisy, entrapment, as well as background relationships between blamer and blamee – for instance, being a romantic partner to a person who is guilty of infidelity. We saw how some of these facts, like hypocrisy and entrapment, seem to have a negative effect on standing, whereas others, like being a partner, seem to have a positive effect. But what is that effect?

Binary is a particular type of response to this question. The following passage by James Edwards (2018) provides a useful example of such a response (it does not literally imply a commitment to Binary, but it does illustrate how someone who does hold this view thinks about standing-relevant facts):

Sometimes facts about A *ground* A's standing to hold B responsible. Such facts give A standing that A would not otherwise have [...] Facts about A may, however, also play a second role. Those facts may *defeat* A's standing to hold B responsible. When such facts obtain, A loses standing A would have otherwise had (Edwards, 2018, p. 7, my emphasis).⁴

The negative effect of standing-relevant facts, according to Binary, is a *disabling* effect – negative facts ‘defeat’ standing; they make us ‘lose’ standing, or take it away. Correlatively, any positive effect of standing-relevant facts is an *enabling* effect – these facts ‘give us standing’ where we otherwise would not have it. What Binary says is that these two effects are the *only* possible ones – either facts make us have standing, or they make us not have it.

⁴ Also compare Isserow's & Klein's (2017)'s remark:

When we declare that we are “not in a position” to criticize others, we refer to some kind of “disabling fact” about ourselves that undermines the illocutionary force of our utterance (2017, p. 198).

Now note the following implication of this view: the issue of standing is overdetermined where there are multiple positive or negative facts. That is to say, where some negative fact disables standing, that is all there is to say about it – it’s gone! No further negative fact has a negative effect, and no positive fact has a positive effect. Similarly, where a positive fact enables standing, that is all there is to say – it’s granted! Unless a negative fact is also present and disables standing, no further facts have any effect on standing. So, we can summarise the view as follows:

Binary: Either standing-relevant facts disable or enable standing altogether, or these facts have no effect on standing.

Obviously, if Binary is true, then standing has to be the sort of moral property that can be enabled or disabled. Let us now look at three reasons why we might think that standing is such a property.

i. Permission?

One binary category is a moral permission. Perhaps certain kinds of conduct are closer or further away from being permissible or impermissible. But, something is either permissible or not; conduct cannot be more or less permissible or impermissible. It might seem natural to think of standing as a permission to hold others accountable. As Wallace (2010) points out, the idea that someone is ‘engaged in an activity that they lack the standing to engage in’ is a ‘fundamental moral objection’ to what they do (2010, p. 317).⁵ One possible moral objection to engagement in some activity, of course, is that the person engages in it *wrongly*. Indeed, the argument in the first half of this thesis was that this is true of hypocritical blaming: when Blake blames Amy for similar wrongdoing, and fails to acknowledge his own wrongdoing, he not

⁵ Also compare Herstein (2020), who holds the view that the lack of standing makes ‘otherwise permissible [blaming] interventions impermissible’ (2020, p. 6). Edwards (2018) also claims that having a duty not to blame (i.e. lacking a permission to blame) is sufficient to undermine standing (see note 11 below).

only has defective standing to blame her, but also blames her wrongly. However, this is not in fact a necessary feature of defective standing; it turns out to be a coincidental feature of that particular case. Defective standing and wrongful blaming are not the same thing, nor are they coextensive.

Here is a general reason not to equate defective standing with wrongful blaming: standing considerations, it seems, play a role in *grounding* wrongness considerations. The fact that a person has defective standing, in other words, seems to play a role in explaining why a particular act of blaming is wrong. We can contrast this with the idea that standing has an *upshot* role. On an upshot view, defective standing is just a consequence of the fact that blaming is wrong. However, the idea that standing grounds wrongness is more intuitive than the idea that it is an upshot of wrongness. Imagine, for instance, that Blake asks Amy, ‘What’s wrong with my blaming you?’. She might answer, ‘You lack standing to!’ That answer is meant to be informative; it tries to offer an argument for why Blake’s blaming her is wrong. But that would not be true if defective standing was just an upshot of the fact that he blames wrongly. Amy’s response would not give him new information. It would only repeat the answer for which he seeks an explanation.⁶

Standing, furthermore, is quite clearly a separate issue from whether blaming is permissible overall. In Chapter 5, we observed that standing is not determined solely by the consequences of blaming; yet the consequences of blaming can be decisive in determining whether blaming is permissible. Recall how it might be impermissible for me to blame my colleague for insulting me when I could instead be treating a patient: this fact has no relevance to whether I have good standing to hold my colleague accountable. Similarly, as we observed in Chapter 4, it may be permissible to blame someone hypocritically given the costs of non-

⁶ There is one way in which what Amy says may be informative: it might be that Blake’s defective standing constitutes a *certain kind* of wrongful blaming. So, in telling Blake that he has defective standing, Amy might be saying that he is committing a particular type of wrong – a *standing wrong*, if you like – rather than simply saying that he’s acting wrongly. But, as we’ll see below, having defective standing does not make it *pro tanto* morally wrong to blame.

hypocritical blaming. For example, it may be permissible for a mob-member to blame hypocritically if acknowledging wrongdoing risks his getting shot. But that, I take it, does not make a difference to the mob-member's defective standing.

Thus, a simple equation of standing with a permission to blame is false; standing is not an all-things-considered permission to blame. A better view is much narrower: standing, some might suggest, is a moral permission to blame in virtue of *only and all of the standing-relevant facts* that are in play in a given instance of blaming. On this view, where the standing-relevant facts in a given instance of blaming do not make blaming impermissible, a person has standing, and where those facts do make blaming impermissible, a person lacks standing. This narrower view allows that we might have standing to impermissibly blame someone rather than treat a patient: here the blaming is wrong because of the duty to treat the patient, rather than because of any standing-relevant fact. It also allows that the hypocritical mob-member has defective standing despite being permitted to blame hypocritically. In that case, if we were to consider only the standing-relevant facts – i.e. that the mob-member blames hypocritically – we would conclude that he blames wrongly.

But this view does not work either. To see this, consider a further fact that I gestured at in the last chapter, which is very plausibly relevant to standing: namely, whether the blamer is a *victim* of the blamed conduct. This is another example of what I called in the last chapter a *connective fact* – a fact which connects the blamer to the blamee or the blamed conduct in a certain way. A victim's connection is that she is among the people wronged by the blamee's conduct. Procedural Value can explain why that connection matters to standing, as follows. The value of a victim blaming a wrongdoer is heightened relative to a non-victim, because an accountability between a victim and wrongdoer offers the wrongdoer an opportunity to deliver corrective responses she owes, to the *very person* to whom she owes them. That value is procedural. Recall Radzik's way (2011) of describing the ideal form of a victim-wrongdoer interaction, quoted in Chapter 2: victim-wrongdoer accountability involves a special 'moment of recognition', marked by a '[Wrongdoer]-sees-[Victim]-seeing-[Wrongdoer] dynamic' (2011, p. 596). Beyond being instrumentally valuable, this dynamic seems valuable in itself, for both

wrongdoer and victim, given that it is valuable in itself that one party complies with her corrective duties, and that the other party gets her due. But, again, we need not rely solely on Procedural Value here. I take it to be uncontroversial that victims have special standing, at least among those who believe in the significance of standing.

Once we appreciate the significance of victimhood, however, we can see that even the narrow Permission view of standing cannot be right. If victims have standing, then *non*-victims lack standing, other things equal: after all, non-victims lack whatever it is that gives victims standing. But, surely, not being a victim of wrongdoing does not make it even presumptively morally wrong to blame others. Non-victims hold wrongdoers accountable all the time, whether on behalf of victims or others, and other things equal, we tend to think this is not in any way wrong, at least when they blame the person accurately and proportionately. So, whatever it is that Wallace refers to as the ‘fundamental moral objection’ to standing-defective blame, it cannot be that there is something *pro tanto* morally wrong about such blame. Hypocrisy is just a special case where standing-relevant facts often play a part in making blame wrong.⁷

ii. Authority?

Another common view is that standing is a kind of authority. Some may think that a person can coherently have more or less authority.⁸ Still, a person either has some authority or none

⁷ Even in the hypocrisy case, standing-relevant facts – i.e. that the person rules out the value of mutual deliberation – is only part of the explanation for why blaming is wrong; it is not sufficient to make blaming wrong, as the mob-member case illustrates. This is due to something we saw in Chapter 4: generally, failing to achieve a certain kind of value is not in itself impermissible. For example, we also need to know what the costs are of achieving the relevant value.

⁸ Not all views of authority can countenance this. Raz’s view (1986), for example, is that your having authority means that others are better off acting on your say-so than acting on their own judgement. Either this is true or it is not. Acting on one person’s say-so might make others even better off compared

at all, and a prominent binary view of standing is that those who have standing have a particular type of authority, whereas those whose standing is defective have no such authority at all. But I do not think that blame involves an authority. And even if it does, it is hard to see why losses in standing generally take this authority away.

The sense of authority specifically associated with standing is the power to place someone else under an obligation (or more moderately, the power to create a moral reason for someone else) through our acts. Again, this appeals to a particular view of blaming – some think what a blaming interaction typically tries to do is obligate someone to respond to wrongdoing in a particular way, *by demanding* that they respond in that way.⁹ We might think standing is a success condition for that function: those with standing can, through demands they communicate in blame, obligate (or create reasons for) the blamee to respond, whereas those without standing lack this power.¹⁰

Edwards (2018) and Ori Herstein (2017; 2020) argue this in some detail. They differ on exactly what reasons or duties a blamer with standing might create. Edwards, for example, suggests that standing is a power to make the blamee duty-bound to offer a reply to the content

to another person, but perhaps the natural thing to say here is that the latter's authority is more important, rather than that they have more of it.

⁹ A different authority view, which I also reject, is that hypocrites lack authority and standing because others have reasons to question the truth of their criticism, given that they fail to 'walk the talk' (Isserow & Klein, 2017). Generally, I doubt that whether a person 'walks the talk' is a reason to question the truth of their criticism, or indeed that this person is knowledgeable about morality. For more on this, see Chapters 1 and 3.

¹⁰ A still different binary thought is that a hypocritical blamer literally cannot *blame*, let alone blame successfully, because a person literally cannot blame without taking its content seriously (Riedener, 2019). I find this is difficult to believe. As King notes, 'hypocritical and meddling blamers, even if they lack standing, do not *fail* to blame. It's precisely their blame to which we're objecting' (2020, p. 2070). A related idea is that blaming is *pointless* where the blamer is not committed to its content (Shoemaker & Vargas, 2019). I suspect both these views are vulnerable to a version of what I say below.

of the blame. On this version of the Authority view, if Blake has standing, Amy has a duty, *because* Blake blames her, to offer a reply to the content of the blame, and if Blake's standing is defective, Amy has no such duty.¹¹ Herstein defends a stronger claim. If Blake's standing is defective, he argues, Amy is permitted not merely to fail to reply to the content of his blame, but also to exclude that content from her deliberations about how she ought to behave. If Blake has standing, in contrast, he can give Amy reasons to focus on, rather than ignore, what he says because of his say-so.¹²

This is a plausible analysis of people's typical responses to defective blaming. As Gerald Cohen's early musings on standing illustrate (Cohen, 2013), it is natural to try to deny

¹¹ Edwards adds a further condition: beyond the ability to obligate the blamee, the blamer must also have an 'agent-relative privilege' to hold the blamee accountable – i.e. must lack a duty *not* to hold the blamee accountable in light of facts particular to herself (2018, p. 17). But this is not fully argued for. Edwards uses the idea of an agent-relative privilege to distinguish his own view from the simple Permission view discussed above: we may lack a privilege (i.e. have a duty not) to hold someone accountable given facts about ourselves, he points out, without it being all-things-considered impermissible to do this (p. 18). Edwards also distinguishes between having a normative power and having a privilege. He uses the example of a referee giving a mistaken penalty to show that a player can have a duty to obey the referee's decision even though the referee has a duty not to apply the rules incorrectly (p. 19). However, this tripartite distinction – between privileges, all-things-considered permissions and powers – gives no particular reason to think that an agent-centred privilege to blame is a necessary condition on standing. For example, my sense is that a referee has standing to give a mistaken penalty. Similarly, as I have argued in the previous chapter, a person who lacks sufficient evidence about someone else's wrongdoing may lack a privilege to blame in light of facts particular to herself. But this person (for instance, an unknown victim of infidelity) might have standing.

¹² Others frame the problem in a similar way – a common thought is that blamee can dismiss blame if the blamer's standing is defective. See e.g. PEA Soup (2019); Lippert-Rasmussen (forthcoming; 2018; 2013); Todd (2017); Friedman (2013, p. 281); Radzik (2012, p. 177). Sliwa (forthcoming) also believes that blame 'alters the normative landscape'.

a critic's authority in the ways mentioned above: by trying to silence them, for example, or to ignore or dismiss their criticism and walk away. However, this natural response is not right.

Let us think again about the case of victims. This may seem the most plausible example of the authority view: as noted, when a victim holds a wrongdoer accountable, she initiates a process in which the very claimant of the wrongdoer's corrective duties demands compliance with those duties. It seems disrespectful not to reply to a victim of ours who is presently blaming us, or to ignore what she says, or walk away; perhaps more clearly disrespectful than it would be to ignore a victim of ours whom we bump into on the street, for example, and who does not hold us accountable. It is tempting to think this is why victims have special standing – because they have a special normative power to obligate the blamee to respond.

However, that is not the only possible explanation of why wrongdoers must respond to victims. Nor is it the most plausible. To understand whether victims have a normative power to obligate *through blame*, what we need to know is not whether victims are owed responses by wrongdoers, but whether *their blaming* wrongdoers is an additional source of wrongdoers' corrective demands. An alternative view is that we owe victims duties of respect in general, and the context affects how these more general duties ought to be discharged; and, accurate blaming is a context where giving a response to blame is a particularly appropriate moment to discharge certain corrective duties. It is hard to see what we gain by introducing the further thought that a victim's blame creates new reasons or duties. What grounds a blamee's reasons to respond to their wrongdoing when blamed for it, is not the fact that someone is blaming them. What grounds those reasons is what independently grounds wrongdoers' corrective duties: viz., facts about their wrongs, and the demands those wrongs create.

To be sure, a victim may have *some* special normative powers. For example, victims seem to have a power to forgive wrongdoers, which is sometimes understood as a power to

release a wrongdoer from certain corrective duties.¹³ But this, notice, is not the same as the power at issue: the power to *release* a wrongdoer from a duty by forgiving them, is different from the power to *create* new duties by holding them accountable. And the latter is more mysterious than the former.¹⁴ It is easier to see the difference made by forgiveness. Imagine two contexts where Amy is unwilling to open herself to an accountability conversation with Ciara for abusing her. In one context, Ciara had sincerely communicated forgiveness to Amy; in another, she had not. Here it is somewhat intuitive that Amy violates a more stringent duty in the first context compared to the second. In contrast, it is harder to see the difference made by blaming. Again, assume Amy is unwilling to be accountable to Ciara. In one context, Ciara has engaged Amy in an accountability process; in another, Ciara has not. Here it is not intuitive that Amy violates a more stringent duty in the second context compared to the first.¹⁵

Some will no doubt demur. It may seem that although a failure to respond to wrongdoing is already morally objectionable, it is even more objectionable to fail to respond to wrongdoing when someone holds you accountable compared to when someone does not. After all, when a wrongdoer is being blamed, it is very clear what she ought to do at that particular moment – the blame brings the subject of wrongdoing emphatically to light, and offers an obvious opportunity to reflect on that subject, explain herself, apologise and so on. Failing to respond at *that* point is especially bad. However, this does not show that being blamed by someone affects the stringency of our duties. We can instead explain this by appeal to the fact that being blamed, especially by a victim, is a particularly apt occasion on which to discharge duties of a given stringency, and it is more blameworthy not to comply with our duties on apt occasions. Thus, rather than a duty- or reason-creating mechanism in itself, blaming may be a mechanism for encouraging compliance with duties or reasons we already

¹³ See e.g. Bennett (2018) Warmke (2016) and Owens (2012, ch. 2). Note that Bennett (2018) also thinks that forgiveness creates new duties in the victim.

¹⁴ Compare Chang's discussion of the more controversial idea that normative powers are generally mysterious (Chang, 2020).

¹⁵ This provides an argument for Matt King's intuitions (King, 2019, pp. 274-275)).

have, or a particularly good context in which so to comply.¹⁶ That helps explain the appeal of wrongdoers' needs to respond to blame, whilst avoiding the mysterious thought that blaming creates new reasons or duties. It also gives a reason to deny the Authority view more generally: if it is doubtful that *victims* have normative powers in blaming, it is even more doubtful that the whole set of people with standing have such powers.

What is more, even if we deny the above – even if we see blaming as a normative power – standing does not plausibly depend on whether the relevant power is enjoyed. One way to see this is to consider the example, sketched in the previous chapter, of a *hypocritical victim*: when a victim blames a wrongdoer hypocritically, his standing is defective, yet it is not plausible at all that the wrongdoer has no duty to respond.¹⁷ Suppose, for example, that Blake in *Supremacists* is a victim of Amy's racial abuse. Suppose he fails to acknowledge his casteist abuse whilst blaming Amy. Blake's standing to blame here is surely defective. The fact that he is a victim affects neither his hypocrisy nor its objectionableness; so whatever negative effect hypocrisy has on standing, that effect is occurring here. However, it hardly follows from Blake's defective standing that Amy may permissibly fail to reply, or walk away, or exclude the content of what Blake says from her deliberations. If victims do have a normative power to obligate wrongdoers, this must ultimately depend on the fact that they are wronged, and are thus owed a response. Hypocrisy does not negate those wrongs, or the duties to respond. Only an unduly cruel interpersonal morality would hold that the strength of our claims on those who wrong us to apologise, compensate, etc., depend on whether we are willing to acknowledge having done something similarly wrong. Just as I should compensate you for stealing your TV even though you stole someone else's TV and did not acknowledge this when you blamed me, I should also apologise and explain myself to you, even though you are not apologising to your victim. So, I think it is both inapt and wrong for the blamee either not to

¹⁶ This echoes problems with the more general idea of second-personal reason-giving as conceived by Stephen Darwall (2006): see David Enoch (2014; 2011).

¹⁷ I discuss a particular variant of hypocritical victimhood in the next and final chapter.

reply, or to exclude from deliberation, accurate blame by a hypocritical victim, even though such a blamer's standing is defective.¹⁸ Since a blamer can have defective standing where the blamee has every reason to respond appropriately to blame, lacking a normative power is not the same as defective standing.¹⁹

iii. Entitlement?

A third binary version of standing is that it is a moral right, or entitlement. Like the concept of authority, we might think that some people have stronger entitlements to certain things than others. Still, entitlement is a binary category in the sense that a person is either morally entitled to something or not, and perhaps the commonest binary view about standing is that those who have it have some moral entitlement to blame, whereas those who lack standing

¹⁸ Edwards may respond that the blamer still lacks standing given his privilege condition – i.e. that the blamer has a duty not to hold the blamee accountable (see note 11). But I doubt that a hypocritical victim has a duty not to hold a wrongdoer accountable at all. A victim is permitted to hold the wrongdoer accountable, and if they do, they also have a duty to acknowledge similar wrongdoing. (See Chapter 4, sec. 4).

¹⁹ Perhaps the problem here is not the Authority view in itself, but what particular versions of it hold that the blamee has a duty to do. For example, Kasper Lippert-Rasmussen has suggested to me that rather than a duty to respond to or consider the content of blame, the relevant duty created by an authoritative blamer might be a duty to endure the emotional costs of blame. So, when a blamer has standing, the blamee ought to endure these costs, but not so when the blamer lacks standing. However, I cannot see why a blamee is ever under a duty to endure the emotional costs of blaming so long as the aims of blame are achieved: so long as the blamee absorbs the content of blame and responds appropriately, there is no obligation to endure emotional attacks. If, on the other hand, enduring emotional costs is necessary for a blamee to respond appropriately, then the blamee ought to endure these regardless of who is blaming her. (I also think this is possible even if we are retributivists: even if we think that people deserve gratuitous harm, we may doubt people have a duty to make themselves endure gratuitous harm.)

have no moral entitlement to blame.²⁰ But, once again, blame is not helpfully understood as an entitlement, and in the sense that we do have blaming entitlements, it is doubtful that standing-relevant facts affect them.

Entitlements are not quite the same as permissions. It may be wrong to do things we are entitled to do. I may be entitled not to marry someone for racist reasons, say, or to speak my mind and say offensive things. Similarly, it may be permissible to do things we are unentitled to do. I may permissibly break into your house to escape the blizzard. So, unlike the Permission view, the Entitlement view might make better sense of the difference between defective standing and pro tanto wrongful blaming: a possible view is that the mob-member who risks getting shot is permitted, but not entitled, to blame hypocritically, whereas the doctor who must treat her patient instead is entitled, but not permitted, to blame her colleague.

However, blame and standing are not really a matter of entitlement.²¹ Ordinarily, entitlements, such as the right to choose a marital partner, or to free speech, are grounded in certain interests of a self-regarding nature. We are entitled to these things because they are central to our ability to shape key aspects of our lives. But blaming is not quite like this: blaming is typically about other people's interests – it is about emphasising the nature and gravity of wrongdoing, the claims of victims of wrongdoing, and the duties of wrongdoers. Of course, when victims blame wrongdoers, they voice their own interests and protest against their violation, and they have powerful interests in the ability to voice these facts.²² But even here, the interest that blaming serves is not primarily the victim's interest in being able to shape her own life, free from external control. Rather, the interest is in influencing others' lives in a way that urges them to make things right.

²⁰ See e.g. Fritz and Miller (2018); Todd (2017); Bell (2012); Cohen (2013); Lippert-Rasmussen (2013) and Radzik (2011). As should be clear from these references, some discussions of standing run different binary ideas together, implying that standing is some combination of them.

²¹ Compare King (2019, pp. 270-276).

²² Compare Smith (2013).

Blaming no doubt does serve some entitlement-grounding interests; for example, the interest in engaging others' agency, and in exchanging and expressing moral views – plausibly these things also help us shape our own lives. We might then understand an entitlement to blame as part of our ordinary entitlements to free thought and expression: the ability, through blaming, to express our moral views and assessments may be part of a more general ideal of a self-determining life in which people have certain prerogatives to think and talk independently. The problem, though, is that this general set of entitlements is surely not disabled by *standing-relevant facts*. Compare accurate blaming with the other things I just suggested we may be entitled to do, like saying offensive things, or marrying people for racist reasons – these things are seriously morally bad. They are surely much worse than accurately blaming others hypocritically, for example, or accurately blaming people when it is not our business. Indeed, it is not a stretch to think that the entitlement to free expression even entitles us to blame people without warrant: after all, unwarranted blaming is not obviously morally worse than much free speech we tend to protect. Yet, blaming someone without warrant – someone we have no reason to think guilty, or who is not guilty – is often going to be morally much worse than accurately blaming someone with defective standing. So, if we can be entitled to blame in the way that we are generally entitled to express ourselves, we should doubt that standing-relevant facts take that entitlement away.

Overall, these substantive problems may not be decisive against the three prevailing versions of Binary, and other versions may fare better. But the problems are serious. Indeed, some who agree with what I've just argued may be inclined to doubt that standing is morally significant at all, given that they think of standing as a binary category.²³ However, as I will show in what follows, standing should not be understood in a binary way in the first place. *Scalar* is a more persuasive account of the importance of standing-relevant facts.

²³ This is one way to interpret sceptics about standing like Lang (forthcoming), King (2019) and Bell (2012).

2. Standing as Appropriateness

If standing is not a permission, entitlement, or authority, then what is it? I will suggest that it is one kind of assessment of the moral appropriateness of blaming, and that this view supports Scalar.

Before making this point, let's further clarify Scalar. Recall that in contrast to Binary, Scalar says that standing comes in degrees: rather than being an all or nothing matter, we can have defective standing to blame without losing standing altogether. This also has implications for how standing-relevant facts work. The effect of some negative standing-relevant fact (like hypocrisy or entrapment), need not be to *disable* standing: to take it away, or deprive a person of it altogether. Rather, the effect can be to *decrease*, *weaken* or *diminish* standing: to make standing *deficient*, rather than absent. Similarly, the effect of some positive standing-relevant fact, such as whether the blamer is a partner or a victim, need not be to *enable* standing: to give standing to us where we would not otherwise have it. Rather, the effect can be to *increase* or *enhance* standing. The relevant facts can give us a particularly high degree of something we already have, rather than giving us something we would otherwise lack.

This allows for a possibility which Binary rules out: once a fact has made a difference to a person's standing, another fact can make a further difference without enabling or disabling standing. For example, where some fact makes a positive difference, a further positive fact might come into play and *further* enhance a person's overall standing, taking it higher up the scales. Likewise, a negative fact might come into play and diminish standing, bringing the level down. And vice-versa: if a negative fact is at play, another negative fact can come into play and diminish standing even further, whilst a positive fact can come into play and make the degree of standing higher. Thus, we can state this idea as follows:

Scalar: Standing-relevant facts do not always disable or enable standing entirely. And when they do not disable or enable standing entirely, standing-relevant facts can increase or diminish standing.²⁴

i. Appropriateness: scalar in contribution and in nature

What would standing need to be for it to satisfy this definition? One thing it might be is a type of moral appropriateness in blaming. This is a promising idea, and is an implication of a combination of two things defended in this thesis so far: Chapter 2's account of blame's appropriateness, and Chapter 5's Procedural Value account of which facts are standing-relevant. I argued in Chapter 2 that appropriateness in blaming depends on how well blame achieves the values that are internal to the practice. Notice, moreover, that the values I claimed to be internal are *procedural*: they are about employing a particular kind of process, of discursive engagement, to achieve reform. Furthermore, I argued in Chapter 5 that standing-relevant facts affect the procedural value of accountability: the value, that is, of particular people employing discursive means of getting particular others to reform. It follows from these two ideas that standing correlates with the appropriateness of blaming: the more or less standing a person has, the more or less their blaming is procedurally and internally valuable – hence, the more or less their blaming is apt. But again, that particular argument is not needed here. The view that standing is a type of moral appropriateness is not controversial. Many commentators suggest something like this view in passing.²⁵ However, most have fully failed to appreciate its scalar implications.

²⁴ Notice this allows that a person sometimes has no standing whatsoever. The thought is simply that no standing whatsoever is not a central consequence of defective standing.

²⁵ Examples: Kyle Fritz and Daniel Miller (2018) write:

The fact that [someone] lacks the standing to blame [...] is a consideration that weighs heavily against the appropriateness of [...] blaming [...]. But [...] standing is only one consideration in the ethics of blame.... (2018, p. 119)

Similarly, Kasper Lippert-Rasmussen asks:

To bring out the scalarity in appropriateness, let us bracket standing-relevant facts for a moment, and focus on other facts that seem to make blame morally appropriate or inappropriate. One fact, of course, is whether blame is warranted: if evidence suggests that a person is not blameworthy, or if they are not in fact blameworthy, it is morally inappropriate to blame a person. But, as observed in Chapter 2, a number of other things make a difference to the moral appropriateness of blaming. *Social setting* is one example. Recall the case of blaming someone for a peccadillo at their funeral; it seems morally inappropriate to blame in this way partly because of the social function of that type of event. Similarly, it seems morally appropriate to blame a person in a social setting where victims of their wrongdoing are present. Another factor I hinted at is *delivery*: blaming someone via jokey text message for something seriously wrong does not appropriately communicate the gravity of what they did. Blaming this person in a sombre tone to their face, in contrast, is a morally appropriate instance of delivery. And there are many more such factors. A third example is *uptake*: it seems morally inappropriate to blame someone while they are comatose, for instance, and appropriate to blame someone when fully responsive.

Once we reflect on these broader appropriateness-relevant facts, the idea that appropriateness is scalar becomes compelling – the moral appropriateness of blame is too complex and multi-layered to be an all or nothing matter. Rather than being either altogether inappropriate or altogether appropriate, blaming can be more or less so.

One reason for this is that there are *multiple* appropriateness-relevant facts, each capable of contributing to a particular case. Sometimes this is true of positive facts – different sources of positive appropriateness contribute by *coalescing* to increase overall appropriateness. For example, suppose I blame someone when their victims are present. That seems

[W]hat, exactly, is the nature of the inappropriateness alleged in [standing-related] responses to blame (2013, p 295)?

And Angela Smith includes standing among:

the various considerations which [...] have a bearing on the kinds of attitudes and responses it is appropriate for us to take toward one another (2007, pp. 475-477).

appropriate, but now suppose I also blame this person when that person is fully responsive. This, plausibly, is more appropriate overall than blaming a comatose person in the presence of victims, as it is not only setting-appropriate, but also uptake-appropriate.

At other times, negative appropriateness is what's compounded by different sources. For instance, imagine blaming someone for serious wrongdoing. It seems inappropriate to do this via jokey text, but even more so to send them a jokey text when knowing that their phone is broken. This is more inappropriate overall, because it is not only delivery-inappropriate, but also uptake-inappropriate.

Another contribution of different dimensions of appropriateness is that some can *countervail* or *attenuate* others.²⁶ Sometimes this happens when inappropriateness in one dimension countervails or attenuates appropriateness in another. Recall, for example, the case of blaming someone in the presence of victims: clearly, it is less appropriate to do this in a jokey tone, compared to a sombre tone. Though here the setting is appropriate, the blame is made less appropriate overall by the delivery in that setting (perhaps as the delivery is insensitive to the setting).

Vice-versa, positive appropriateness contributions can also countervail or attenuate negative ones. Suppose I blame someone for a peccadillo at their funeral: it may well be more appropriate to do this in a light-hearted tone, compared to a sombre tone. Though here the setting is inappropriate, the blame is made more appropriate overall through the delivery (again, perhaps as the delivery is sensitive to the setting).

A second way blame can be more or less appropriate is as follows: *each* appropriateness-relevant fact is scalar in nature, in the sense that each can apply more or less powerfully in a given case. For example, while it is setting-appropriate to blame someone in the presence of victims, these victims might be more or less present – they might be listening attentively, or half-listening, or barely within earshot and vaguely listening. Similarly, though it

²⁶ I borrow the language of contributing, countervailing, enabling and attenuating from Dancy's discussion of reasons (2004, ch. 2-4)

is uptake-appropriate to blame someone while they are responsive, the person may be more or less responsive – they might be tipsy, or drunk, or steaming drunk. At each step, the appropriateness of blaming is not ‘switched off’; it only decreases.

And we see a similar pattern when it comes to the negative appropriate-relevant facts: each is scalar in nature. For example, it is delivery-inappropriate to blame someone for serious wrongdoing in a jokey tone. Yet, that jokiness can be more or less unseemly. The blame could be sarcastically jokey, or it could involve stifled laughter, or hysterical laughter. Again, inappropriateness is not ‘switched on’ at each step; it only increases.

This, basically, is how I will argue standing-relevant facts also work. As with other appropriateness-relevant facts, each standing-relevant fact can make a stronger or weaker contribution, and multiple such facts can contribute to a given case, with the result of increasing or decreasing a person’s overall degree of standing. While Scalar may initially be surprising, then, it should not be so surprising having seen that other appropriateness concerns work this way.

ii. Contrastive implications

Before arguing that Scalar is true, we should briefly consider what is at stake in the dispute as I have presented it. Why should we care whether standing is an assessment of moral appropriateness, rather than say moral permissibility, authority or entitlement? It is clear what normatively follows from the idea that standing-defective blame is morally wrong, or unentitled, or unauthoritative. What, though, are the normative implications of the scalar view? My proposal: standing matters in a *contrastive* way. When there are various options in how and by whom blame should be issued, standing considerations guide us on which options to select. They inform us about whether we ought to blame a person rather than allowing someone else to do so; whether we should encourage or discourage blaming by a particular person rather than another person; and indeed whether we should encourage or discourage blaming someone in one way rather than another.

Here is the justification for that view. We have stronger moral reason (all else equal) to instantiate more appropriate blame rather than less appropriate blame. This principle explains why we have stronger moral reason, for example, to blame in the presence of victims, rather than in their absence; why we have stronger moral reason to blame sombrely rather than jocularly in response to serious wrongdoing; and stronger moral reason to blame someone while they are awake rather than while they are comatose. As it is one element of moral appropriateness, I submit a similar principle about standing: we have stronger moral reason (all else equal) to ensure that blame is issued with more standing rather than less.²⁷

The significance of this principle can be seen in ‘piling on’ cases, where we know that if we do nothing, lots of people are going to accurately hold someone accountable, resulting in gratuitous harm, or even counterproductive, harm on the blamee. Here we want some people to accurately blame, but not all, so we need to know whom to select. Scalar implies that in such cases, we ought to encourage those with more rather than less standing. For example, consider the following version of the Estrangement example explored in Chapter 5. Suppose Blake is a stranger to Amy and Ciara, and is considering whether to blame Amy for cheating on Ciara. Blake can see that Djamel, a close friend of theirs, is about to blame Ciara. He can also sense that it would be too intimidating for both himself and Djamel to blame Amy together. On the view I defend, Blake then ought to leave the blaming to Djamel, because Djamel has more standing than himself. And the same applies to third parties. Suppose Amy’s mum is watching on, knows about the cheating, and can see that Blake and Djamel are about to confront Amy. Amy’s mum ought to discourage Blake, because Djamel has more standing.²⁸

²⁷ This point is particularly relevant in legal and political contexts, where we must assign institutional roles to a select group of people to hold particular citizens, such as criminal wrongdoers, to account. The upshot of this discussion is that, all else equal, we should assign this role to those with more rather than less standing. For debates about standing in these contexts, see Beade (2019), Shelby (2016), Watson (2015), Duff (2010), Tadros (2009) and Matravers (2006).

²⁸ ‘Better’ or ‘worse’ standing may be more accurate, but I tend to say ‘more’ or ‘less’ here to underscore the basic point of scalarity.

The idea that blame should be issued with more rather than less standing also has implications for who ought to endure costs in blaming. Personal costs to us are less significant when we know that there will be a loss in standing-appropriateness if we do not bear them. Thus, imagine Amy's mum can either discourage Blake or Djamel to blame Amy for cheating. She knows that discouraging Djamel will be much easier, as she knows him, whereas discouraging Blake will be harder as Blake is a stranger. Amy's mum then ought to bear greater costs to discourage Blake in this case, as not doing so will result in a significant loss in standing-appropriateness.

Likewise, where we *know* that blame is going to be issued with as much standing as is feasible,²⁹ considerations of personal cost make more of a difference to what we ought to do. Suppose Amy engages in infidelity with a colleague in an unprofessional manner. Blake is Amy's new manager, and he has two options. He could blame Amy non-hypocritically, by fully acknowledging the fact that he himself cheated with a colleague; but this would be costly to him, as it would risk professional penalties. But Blake has a co-manager, Djamel, who is also guilty of infidelity with a colleague, and would be willing to blame Amy. Blake also knows that Djamel is willing to acknowledge his infidelity – everyone already knows about Djamel's misdeeds, so *his* acknowledgment will not be professionally risky. If Djamel blaming Amy will be adequately effective, Blake might then be permitted to forgo the relevant costs. He can perhaps persuade Djamel to do the blaming instead of himself, as this creates no overall loss in standing-appropriateness, and is no costlier for Djamel.

3. Standing: Scalar by Contribution and Nature

²⁹ I gloss over the issue of what full or maximal standing would be. On the Binary view, having full standing and having standing are the same thing. But if what I go on to say is right, very few will people ever have full standing. A person with full standing has all the positive facts in her favour, and no negatives against. But not only that: the positive facts also each have a maximal degree of significance. It is tricky to imagine a person who meets all of the above, partly as there are multiple standing-relevant facts, and partly as it is hard to gauge limits on the positive or negative impact of each fact.

So much for what might explain a scalar view, and what its implications are. These implications may be plausible enough, but they do not show that Scalar is true. We might agree that blame should be issued more rather than less appropriately, but deny that this has anything to do with standing. In this section, I offer further argument in support of Scalar. Through careful examination of a core range of standing-relevant facts, we can see how Scalar explains their force in a much more intuitive way than Binary.

i. Scalar by Contribution

Like other appropriateness-relevant facts, standing-relevant facts can contribute to a single instance of blaming to increase or decrease standing.

a. Coalescing

Let's start with positive facts. Recall how blaming someone in the presence of victims *plus* doing this when they are fully responsive makes blame more apt. Analogously, a greater number of positive standing-relevant facts can coalesce to increase a person's standing.

To see this, consider two types of positive facts already mentioned: victimhood and having a certain kind of relationship with the blamee, such as being their partner or close friend. Imagine Amy cheats on her partner, and compare three things that may be true of Blake. First, Blake might be a *distant victim* of Amy's cheating. This might be the case, for instance, where Blake barely knows Amy and Ciara, but he knows that his partner, Djamel, knows Amy and Ciara, and he also knows that Amy's infidelity encouraged Djamel to treat him badly. Blake's victimhood, despite it being of a distant kind, has a positive effect on Blake's standing: his standing is not intuitively the same as someone who is totally unaffected by what Amy does.³⁰ A second possibility is that Blake is a close friend of Amy's and Ciara's. Again this seems to give Blake a special standing to participate in accountability: his standing is not the same as a total stranger, for example. A third possibility is that Blake is Amy's partner. This

³⁰ Compare Hughes & Warmke (2017 sec. 4).

likewise has a positive effect: as we saw in the Estrangement case, a partner's standing is not the same as a stranger's.

Binary holds that if these facts make a difference, they each 'switch standing on'. It is not the case that victims have more standing than non-victims, friends have more standing than strangers, and partners more than friends. Rather, if these different facts have a positive effect, they each make it the case that a person has standing. But this, I will show, is too crude an analysis. All three possible types of fact about Blake have independent force, and when they combine – when more than one is true of Blake – they intuitively increase his standing,

A possible way to support the contrary view is to argue that a common factor in all three cases enables standing: namely, that Blake is a *stakeholder* in what Amy did. The distant victim, the friend and the partner all have a stake in Amy's infidelity. We might think that each of these ways of making someone a stakeholder is like an on-off switch – they are all ways of giving a person standing to hold someone accountable where they would otherwise have none. This is implausible, however, if by 'enabling standing', we mean, 'being a necessary and (potentially) sufficient condition for a person to have standing'. Surely non-stakeholders in wrongdoing have standing to hold wrongdoers accountable – to think otherwise is to deny that anyone has standing to hold anyone else accountable for non-person-affecting wrongs.³¹ But that is not intuitive: some people surely have standing to hold others accountable for committing wrongs that have no stakeholders.

Some may argue that stakeholdership is necessary and sufficient for standing to blame someone for intimate wrongs, such as infidelity: we might think that. only stakeholders in

³¹ This is different from cases discussed in Chapter 2 of wrongs where victim-wrongdoer accountability is not possible, such as cruelty to a non-human animal, or burning, for fun, a few remaining copies of a long-dead author's book. Perhaps these acts have stakeholders. But sometimes we might act wrongly in a way that makes it the case that humans and non-humans are badly off, though without any one of these being a stakeholder in what we do, because none of these beings would exist were we to act differently. No need here to rehearse the immense literature on Derek Parfit's non-identity problem, but see Frick (2017) and Parfit (2017) for recent discussion of various views.

infidelity have standing to blame people for it. But even this idea does not appropriately explain our judgements about standing across the different versions of Blake (as distant victim, close friend and partner). What we want to say is not only that victims, friends and partners all have standing, but that some of these have more than others. In particular, we want to say that the partner has more standing than both the friend and the distant victim.

Why does the partner intuitively have the most standing? Because being a partner combines the importance of two distinct facts – victimhood plus intimate relationships. Blake as the partner, of course, is the primary victim of Amy’s infidelity. But not only that. He is a victim who, like Blake as the friend, *also* bears a historical relationship with the blamee, the value of which, as argued in Chapter 5, depends on accountability practices between them. So, if stakeholderhood does enable standing, it surely enables standing of different degrees, depending on the type of stakeholder a blamer is. Partnership, therefore, provides a clear instance where multiple standing-relevant facts – in this case victimhood-appropriateness and relationship-appropriateness – combine to elevate a person’s standing in the scales. That is why, if we had to select one version of Blake to deliver the blame to Amy, it is clear that we ought to select Blake as the partner.

b. Hypocrisy: disabler or enabler?

We have so far considered the scalar contributions of positive standing-relevant facts. The negative facts are also intuitively scalar in an analogous way. Recall the case of blaming a person for serious wrongdoing via jokey text, plus doing this whilst victims are present; the second factor makes blaming even more inapt. That is true of standing-appropriateness too: different negative standing-relevant facts can combine to further reduce a person’s standing.

Examples of negative facts we have already considered are entrapment – where a blamer lures a person into committing a wrong in order to blame them – and hypocrisy – where a blamer fails to acknowledge guilt of similar wrongdoing. Let us now introduce a third fact that many people take to have a negative effect on a person’s standing: *complicity*. Suppose

Blake eggs Amy on to make an offensive racist remark, before proceeding to hold her accountable for her offense. Clearly his standing to do this is defective.

As above, I claim that each of these factors – entrapment, complicity and hypocrisy – has distinct significance, and so other things equal, a person’s standing decreases every time a new negative fact is introduced. Blake might egg Amy on to make racist remarks, *and* do this in order that he can then hold her accountable, *and* be guilty of a similar act of (casteism) himself. It is powerfully intuitive that his standing reduces at each step.

I earlier sketched what I think is distinctive about entrapping blamers: they execute an intention to practice accountability in a perverse way, by getting someone to commit wrongdoing because this enables them to be warranted in holding them accountable, where the reverse should be true. Ideally, we should hold people accountable because they commit wrongdoing. I also believe that complicity is a distinctive standing-relevant fact. We can again state this in terms of *connective facts*. What connects a complicit blamer to the blamed conduct and the blamee is similar to what connects a hypocritical blamer to these things; however, there is the following difference. Rather than being blameworthy for a *type*-identical act, as is true of mere hypocrisy cases, an accomplice is blameworthy for a *token*-identical act. So, as an accomplice, not only is Blake together with Amy in the need to respond appropriately to wrongs *like* racism; he is also together with her in the need to respond appropriately to the *same* wrongful act of racism. This grounds an even deeper connection between the agents. It is a greater failure of solidarity between them not to recognise that deeper connection. So, other things equal, when a complicit blamer blames, this can involve a greater deficiency in procedural value than when a merely hypocritical blamer blames.³²

How might a proponent of Binary respond? She might mirror the response above, and suggest that while facts like complicity and entrapment may seem to have distinctive importance, they ultimately depend on a common fact: namely, *hypocrisy*. One observation in

³² This is an alternative explanation of Victor Tadros’ and Gary Watson’s intuitions (Tadros, 2009, pp. 398-400; Watson, 2015, pp. 176-184).

support of this idea is that the entrapping and complicit blamer, as well as the merely hypocritical blamer, all have the following thing in common: all are involved in similar wrongdoing. A second supporting observation is as follows. If a complicit and entrapping blamer *fully and openly acknowledges* her complicit and entrapping roles, she seems to restore her standing, and her complicity and entrapment ceases to have a negative effect. We might then think that these are all are species of hypocrisy, and, therefore, that hypocrisy is the common on-off switch in these examples – complicity and entrapment are different ways of ‘flipping the hypocrisy switch’, and disabling standing.

But, even if we agree that hypocrisy is a core aspect of these cases, it does not follow that entrapment and complicity have no independent significance. This comes out forcefully when we hold the fact of hypocrisy constant. Assume, for instance, that Blake will not acknowledge any involvement in similar wrongdoing. One thing he will not acknowledge is his guilt of casteism. That, we saw, makes his standing defective. But now suppose he will also not acknowledge two further facts about himself: not only is he guilty of casteism, he *also* egged Amy on, *and* his intention of getting Amy to act wrongly was to enable himself to hold Amy accountable. We can clearly observe diminishment in standing as these further unacknowledged facts are introduced. This is naturally captured by the idea of an appropriateness scale. Though it is inappropriate to blame whilst failing to acknowledge type-identical wrongdoing, it is even more inappropriate to do this whilst also failing to acknowledge entrapping a person into wrongdoing, and even more so whilst also failing to acknowledge complicity in that wrongdoing. Again, this explains why, if we must select one blamer from this bad bunch, we ought to select the mere hypocrite, or mere entrapper, or mere accomplice, rather than a blamer who has more of these negatives in his disfavour.

In sum, my thought about hypocrisy is similar to my response to the idea that stakeholderhood enables standing – even if hypocrisy enables defective standing, it enables differences in the degree of this defect, for it enables further facts, like complicity and entrapment, to further diminish standing.

c. *Countervailing and attenuating*

Another parallel to consider, between standing on one hand and other appropriateness concerns on the other, is that positive and negative appropriateness facts can countervail or attenuate one another. Again, the same holds for standing-relevant facts.

When discussing other appropriateness facts, I gave the examples of how blaming someone at their funeral in a light-hearted tone countervails the inappropriateness of blaming them, and how blaming someone in a jokey tone attenuates the appropriateness of blaming them when victims are present. For an illustration of how this works for the standing-relevant facts, return to the case of a hypocritical victim. Suppose Blake holds Amy accountable for her racist remark about him, whilst failing to acknowledge that he himself made a casteist remark about Ciara. Blake's hypocrisy, as noted earlier, clearly makes his blame defective. Binary cashes this out with the thought that Blake now has *no standing* to blame Amy at all, despite being a victim. This is not to say, necessarily, that victimhood never has a positive impact. But whatever positive impact victimhood has, Binary implies that it is cancelled in this case, for the negative impact of hypocrisy is to disable standing altogether, and this means cancelling any positive impact victimhood might otherwise have.³³

Yet, this is not a plausible way to understand such cases. Consider again the significance of victims blaming wrongdoers: the prospect of Radzik's 'wrongdoer-sees-victim-seeing-wrongdoer' dynamic, where one person urges another to recognise the implications of what they did to her, and thereby to deliver the responses they owe directly to the claimant.

³³ Like the issue of full or maximal standing, I have ignored the issue whether it is ever plausible that a person has no standing at all, or even negative standing. Again, I don't have a well worked-out conception. Imagine a librarian hypocritically blaming a customer for not returning a library book on time, when they routinely don't return library books on time. Here negatives seem as weighty as positives – we might be indifferent between the librarian blaming and no-one blaming. Perhaps this is zero standing. Now imagine a perpetrator of mass atrocities blaming someone for not returning a library book on time. Here standing-relevant facts seem to make blaming utterly inappropriate – we'd plausibly prefer that no-one blames rather than this person. Perhaps this is negative standing.

All of that is possible even where the blamer fails to acknowledge wrongdoing – as pointed out in Chapter 5, the significance of moral abuses committed against us, as well as the significance of these abuses to our lives, and the value of engaging wrongdoers on these issues, is not undermined by the abuses *we* ourselves commit, or the abuses we acknowledge committing. So, it is very difficult to see how hypocrisy would cancel the positive impact of victimhood.

A better explanation of this case is thus a scalar one. Blake's standing as the victim is not totally lacking; he has *some* standing because he is a victim. Nonetheless, his standing is *deficient*: its degree is lower than it could be – in particular, lower than it would be were Blake to blame Amy non-hypocritically. That is a much more intuitive way to describe the relevance of both the negative and positive facts in this case.

To be clear, this is not the view that standing-relevant facts are *always* additive. I doubt that to determine a person's standing in every case, all we need to do is to add positive facts and subtract negative ones to come up with an overall assessment. As with other appropriateness-affecting facts, there may be *interaction* effects: negative facts might cancel or at least attenuate the contribution of positive facts, and vice-versa. That comes out in the case of blaming someone in the presence of victims – the appropriateness of this is attenuated by blaming in a jokey tone. A credible analogue, when it comes to standing, is a certain kind of *hypocritical partner*. Imagine Amy has cheated on Blake; Blake has also cheated *on her*, and Blake refuses to acknowledge his cheating on Amy when he blames her for cheating on him. As we saw, Blake's partnership with Amy would normally have a positive impact on his standing. Yet this impact, unlike in the hypocritical victim case above, seems diluted by his hypocrisy. One explanation of this is the following. The special appropriateness of partners holding each other accountable, I've been arguing, is grounded in the fact that certain accountability practices constitute the value of partnerships. Yet, when partners wrongly cheat on one another – and all the more so when they are unwilling to acknowledge this – the value of that partnership is impaired, so that, over time, accountability no longer instantiates the value that it once did.

The positive contribution made by this value thus seems at least attenuated, if not disabled altogether.³⁴

But, the point of the hypocritical victim is in keeping with the theme of the discussion so far. Sometimes different standing-relevant facts do not cancel or disable the relevance of other standing-relevant facts; and, where this is true, an intuitive way to describe the distinct impact of each fact is that they enhance or diminish someone's overall degree of standing.

ii. Scalar in nature

One response to all of this is a bit like Todd's idea explored in the previous chapter. I may be right that there are multiple distinctive standing-relevant facts; however, I may be wrong in thinking that these are best-understood as contributing to some *common* assessment of a person's standing. An alternative view is that there a number of radically different types of standing – 'partnership-standing', 'victimhood-standing', 'hypocrisy-standing', let's say – each of which is not commensurable with any others, and each of which is binary. For example, a person might have a 'partnership-standing' switch, or a 'hypocrisy switch', which is either on or off, and in each case, a different type of standing is either enabled or disabled. However, this view is also not intuitive. As I noted about inappropriateness and appropriateness in blaming, *each* appropriateness-relevant fact can have a different degree of significance. The same is plausible for standing-relevant facts, because *each* such fact is scalar by nature. Therefore, standing is scalar even if there are different kinds of standing. So I will now explain.

Consider, for example, how the significance of a victim's presence can be scalar: there is greater appropriateness in blaming a person in the presence of victims when victims are listening attentively, compared to blaming in the presence of victims when they are vaguely listening. The same is plausible for positive standing-relevant facts. For instance, the significance of a partnership to standing increases or decreases given the scalar nature of the

³⁴ Again, more on impairment in partnerships and other wrongdoer-victim relationships to follow in the next chapter.

status of such relationships. Thus, a partner intuitively has more standing to blame another for infidelity when their partnership has been going strong, compared to when it has long been declining. Similarly, the significance of victimhood is a function of the gravity of wrongdoing, which is also scalar. For instance, gravity of wrongdoing is affected by degrees of negligence. Thus, someone hit by a driver who is miles over the alcohol limit has more standing to blame the driver than someone hit by a driver who is one iota over the alcohol limit.

Matters are likewise for the negative facts. When it came to appropriateness-relevant facts, we observed that blaming can be more or less inappropriate in delivery: for example, blaming someone for serious wrongdoing whilst hysterically laughing is more inappropriate than blaming them in a sarcastically jokey tone. Similarly, we can be more or less complicit, entrapping or hypocritical in blaming a person, and the more of each thing we are, the less standing we intuitively have. For instance, as an accomplice, Blake has less standing to blame Amy for racism when he plans Amy's racist remarks together with her, compared to when he merely fails to discourage her from making these remarks when he easily could. As an entrapper, Blake has less standing when he entraps Amy into wrongdoing purely in order to humiliate her through blame, compared to when he entraps her because he deems this the least harmful way to hold her accountable. And, as a mere hypocrite, Blake has less standing to blame Amy for cheating on Ciara when he fails to acknowledge serially cheating on Djamel, compared to when he fails to acknowledge cheating on Djamel just the once.

Furthermore, as we saw earlier in the thesis, a further determinant of how hypocritical someone is, is how appropriately acknowledging of wrongdoing that person is, and this variable also makes an intuitive difference to degrees of standing. For example, Blake has less standing to blame when he deliberately diverts discussion away from his casteism, compared to when he fails to raise the subject because he has forgotten about it. Similarly, Blake has less standing if he switches topic as soon as his casteism comes up, compared to when he diverts the conversation having listened carefully to what Amy has to say about it.

Overall, the scalar nature of standing-relevant facts shows that Scalar is true regardless of whether there is a common measure of standing to which different types of fact contribute.

4. Threshold

That concludes the argument for Scalar as grounded in a type of appropriateness that depends on connective facts about the blamer. Standing is scalar, I have said, not only because different types of fact can together affect its degree, but also because each type of fact can have a scalar effect. Let us finally consider a version of Binary that might respond to these observations.

This view – I call it *Threshold* – accepts much of what I have said above about how standing-relevant *facts* work. It allows that these facts together have a bigger or smaller effect on standing, for instance, by diminishing, enhancing, countervailing and attenuating one another. It also allows that each standing-relevant fact has a bigger or smaller positive or negative effect in a given case, depending on the context and the other facts at play.

However, the thing these facts ultimately determine – viz. *standing* – is still an all or nothing category. Standing depends on the binary question of whether the impact of various facts crosses a certain kind of threshold. On this view, all of the relevant facts in a given case either will or will not imply that a blamer blames wrongly, let's say, or without authority, or without entitlement. If the relevant facts do not have any of these implications, then a person has standing, and if they do, the person has no standing at all. Here is a more technical statement:

Threshold: Beyond enabling or disabling standing altogether, standing-relevant facts can also make a scalar difference to whether standing is altogether enabled or disabled.

However, Threshold remains importantly limited. It overlooks the fact that we have scalar intuitions, not only about standing-relevant facts, but also about how these facts ground certain normative conclusions. Because of this, Threshold cannot appeal to what seems to be the right sorts of reasons to justify certain decisions about how and by whom blame should be issued.

To see the problem, imagine these two quandaries for Amy's mum:

Two Hypocrites: Knowing that Amy will react counterproductively, Amy's mum will refrain from blaming her for infidelity. She can persuade either Blake or Ciara to blame Amy instead.

To avoid ‘piling on’, Amy’s mum will persuade only one to blame. She knows Blake and Ciara, and knows though both will blame Amy with adequate success, they will also blame her hypocritically: both have also cheated on their partners, and are unwilling to acknowledge having done so. However, Ciara is a total stranger to Amy, whereas Blake is her partner (who cheated on his ex³⁵).

Two Intimates: As in *Two Hypocrites*, except the following. Amy’s mum knows that Blake and Ciara would both blame Amy *non*-hypocritically if persuaded to blame. Moreover, both are close friends of Amy’s. However, Blake is also Amy’s partner.

In both cases, Amy’s mum ought to select Blake rather than Ciara. Threshold views cannot explain this in the right way. On Threshold views, as far as questions of *standing* are concerned, Amy’s mum ought to be indifferent about whom she persuades: in both cases, the relevant facts about Blake and Ciara put them on the same side of the Threshold. So there are no reasons of standing to pick one over the another.

Consider *Two Hypocrites*, for example. A Threshold view must say that both Ciara and Blake are ruled-out when it comes to standing, as their hypocrisy will, on any Threshold view, make them blame wrongly, or without entitlement or authority. No reasons of standing, then, to choose one or the other.

Likewise for *Two Intimates*: a Threshold view must say that both blamers are ruled-in, standing-wise, as no such view would say that either of them blames wrongly, or without entitlement or authority. Thus, in both cases, the two are, standing-wise, on a par. Again – no reasons of standing to pick between them.

To be sure, Threshold views are compatible with the *conclusion* that Amy’s mum ought to select Blake rather than Ciara; after all, reasons other than those of standing could explain why she ought to do this. But that is not the point. The point is that, intuitively, standing plays a role in explaining that preference. The *reason why* Amy’s mum ought to choose Blake rather

³⁵ This clause is meant to bracket the complication raised above, that cheating on a person might impair the value of a relationship constituted by accountability with that person.

than Ciara is not a reason independent of standing; rather, it is precisely that Blake has more standing than Ciara. Scalar has the advantage of delivering this more intuitive explanation.

Combined with the problems I raised earlier for thinking about blame and standing in terms of permissions, authorities and entitlements, I think this explanatory advantage is a strong reason to prefer Scalar over Threshold.

Conclusion

Matt King (2019) provides a fitting example of the type of reasoning this chapter cautions against. '[W]hilst hypocritical and meddling blame is inappropriate', King says, 'this fact is not best explained by the notion of standing.' (2019, p. 266) This, I have argued, is not right. It *is* right if we think of standing as we ordinarily do: as a binary category – 'a bar' to accountability (Scanlon, 2008, p. 175), or at least to 'legitimate' accountability (Duff, 2018, p. 781). But it is not right if we think that standing just *is* another element of appropriateness. If that is the better view of standing, then considerations of standing are well-suited to explain the inappropriateness of particular cases. My claim is that this is the better view: like hypocrisy, we should stop conceiving of standing as a bar to blaming, and start conceiving of it as an assessment of the moral quality of blaming.

Perhaps this raises the suspicion that Scalar is not an account of *standing*, but rather of something else. A sceptic about standing may suggest that although Scalar may show that certain facts about a blamer are morally important, it does not show that standing is the morally important category we thought it was, in the sense of showing that it is e.g. a permission, right, or authority. However, I have argued that although Scalar is a foreign idea in general discussions about defective blame, it is less foreign upon a closer look at particular cases. In particular cases, talk of standing in degrees is highly intuitive. While people *think* they think about standing in a binary way, then, perhaps they do not *really* think about standing in that way. At least not all the time.

But at any rate, the sceptical view which accepts Scalar, and also rejects the significance of standing as we ordinarily think of it, is not of concern to me. One way to understand this

chapter is as an account of the moral importance of standing-relevant facts. We can call that importance *standing*, or we can call it something else. The type of scepticism that concerns me denies that standing-relevant facts – hypocrisy, estrangement, and so on – make a moral difference to the ethics of blaming. Scalar addresses that particular scepticism. It shows that in order to make a moral difference, standing-relevant facts need not be permission-, authority-, or entitlement-affecting. Their main importance instead lies in helping us select whom to encourage to blame, how we should encourage them to blame, and at what cost. Rather than rejecting versions of Binary, then, sceptics about the importance of standing-relevant facts should shift attention to arguing against Scalar.

The next and final chapter extends this analysis to one type of hypocrisy that has attracted particular curiosity: namely, where blamer and blamee commit similar wrongs against each other. Applying a scalar view to these cases shows what makes them puzzling. I will address this puzzle, and in doing so, begin to develop a more detailed account of the significance of moral and personal relationships to our accountability practices.

7. Standing, Impairment, and Repair.

This chapter explores one way in which my accounts of hypocrisy and standing might illuminate wider problems in the ethics of blaming. The account of hypocrisy, to recap, is that even if silence can be worse, hypocritical blaming is wrongly second-best given the values of Mutuality that non-hypocritical blaming can achieve. The account of standing was that this second-best-ness also explains the hypocrite's defective standing to blame: the procedure of hypocritical blaming is not as good as non-hypocritical blaming, as the best procedure involves similar wrongdoers reasoning together about their wrongs. However, this does not mean that the hypocritical blamer has no standing at all. As standing is scalar, rather than binary, other standing-relevant facts may countervail or attenuate the negative effect of hypocrisy. So the most we can conclude about the standing of a hypocrite is that it is diminished in the light of his hypocrisy, or that it is deficient overall.

Some philosophers have tended to focus on instances of blaming when hypocrisy combines with one set of facts in particular.¹ In these instances, not only does the blamer commit wrongdoing that is similar to the blamed conduct. He also commits that similar wrong *against the blamee*, and the blamee commits her similar wrong *against him*. This context – I will call it the context of *mutual wronging* – is another plausible instance of standing-deficient blame: I have deficient standing to blame *you* for wronging *me*, if, after all, *I* similarly wronged *you*. But why have these cases attracted special attention as examples of defective standing?

¹ See e.g. Shoemaker & Vargas (2019); Edwards (2018); Lippert-Rasmussen (forthcoming a); Wallace (2010), Scanlon (2008) and Duff (2010).

At first glance, and if my arguments succeed until now, defective standing in these cases may seem easy to explain: it is hypocritical to blame others for wrongs similar to our own without acknowledging our similar guilt, and hypocrisy is a detriment to standing; so the mutual wrongdoer loses standing because of his hypocrisy. However, this response raises a problem. Though I blame you hypocritically when I commit a similar wrong against you, it is also true that you committed a similar wrong against me. This means I'm a victim of what you did. Normally, victimhood has a positive effect on standing, as became clear in the last two chapters. Yet, it doesn't seem to have this effect here: intuitively, a hypocrite who is not a victim has no better standing than a hypocritical victim who is also a mutual wrongdoer. This suggests that hypocrisy does not fully capture why standing is defective in these cases. Something else is at play.

Having explained this issue in more detail in Section 1, I develop an account of what that further fact is. The core reason why the blamer's standing is defective in mutual wronging cases has to do with the blamer's relationship with the person he wronged, and who also wronged him. His loss of standing owes to the fact that he fails to advance prospects of a morally healthy relationship between himself and his blamee. This argument does appeal to the factor of hypocrisy, but not exclusively. It shows how a background of mutual wronging can aggravate the fault of hypocrisy in a distinctive way.

Whilst it is common to think of wronging a person as having implications for relationships, it is not clear how to understand that idea: we don't seem to have anything like a relationship with many people we can wrong. In Section 2, I explain what I think the central relational effect of wronging is. In contrast with other views, such as T. M. Scanlon's, I do not rely on the idea that mutual wronging impairs a *pre*-existing relationship between blamer and blamee. Rather, wronging creates relationship demands: whether you had a relationship with someone or not, wronging a person demands that you interact with them in certain ways. Any future relationship a wrongdoer and victim might have is impaired if these interactive demands aren't met.

These demands are also what explains the mutual wrongdoer's defective standing. In an accountability process, it is appropriate for both blamer and blamee to mutually show recognition of having wronged one another, and a blamer diminishes his standing by impeding this collective reparative project. I argue this in Section 3, and show how it improves on alternative accounts, such as Scanlon's (2008) and Antony Duff's (2010), both of which ground lack of standing in a pre-existing relationship that has been impaired. The alternatives go wrong in equating the significance of purely moral relations, i.e. as agents who might wrong one another, with the significance of relationships in the more ordinary sense, such as collegiate and intimate relationships. The alternatives *do* capture something true about the way in which standing is affected by relationships in the ordinary sense: the value of some relationships depend on historical facts, such as repeated interaction and contingent connections, and their relevance to standing is sensitive to changes in such facts. However, they overlook the fact that accountability works differently between agents who do not share the relevant historical connections. The importance of opening ourselves to accountability for having wronged other agents, qua agents, is not sensitive to the past, and accountability processes are a central way of appreciating this fact.

1. Why not hypocrisy?

Before getting to my account, let me spell out the problem of standing raised by mutual wronging cases, in which you wrong me now, when I have similarly wronged you in the past, and I blame you for wronging me now without acknowledging having wronged you in the past. Why have these cases been used as a paradigm in discussions of hypocrisy? One reason might be that hypocrisy of this kind is especially familiar in daily life. Often the people we hold accountable are close to us; we tend to share more time and space with these people. Sharing more time and space with people can raise the likelihood that they wrong us, and, it can raise the likelihood that, when they do wrong us, and when we have committed similar wrongs, these wrongs were committed against those same people. Example: If I steal my housemate's beer from the fridge, and my housemate is herself guilty of a similar theft, it may be more likely

that her similarly wrongful act is committed against me rather than against a stranger, since both our stuff, rather than a stranger's, is more likely to be in the same place. However, these cases have more to them than simply serving as familiar examples of hypocrisy. Indeed, once we focus on the question of what diminishes standing in these cases, it becomes clear that hypocrisy is not the only thing going on.

i. Mutual and non-mutual cases

To see the issue, let us consider a concrete example:

Co-Battery: Blake, a total stranger to Amy, attacks her just for fun, injuring her. Later, Amy's injury heals and Amy attacks Blake just for fun. The fun has nothing to do with exacting revenge; at the time of the attack, Amy had forgotten about Blake's attack. She injures Blake's arm as badly as he injured her own arm. Blake blames Amy for her wrongful attack.

Other things are equal.

Blake clearly blames Amy with deficient standing. And it is also clear that if he fails to acknowledge his similar wrong against Amy, he blames her hypocritically. It might seem, then, that this is all there is to say about Blake's standing to blame in this case: as is true of other cases we have considered, Blake's standing is diminished because of his hypocrisy.

Consider two reasons in support of this response. The first is to note that there is no difference in Blake's degree of hypocrisy in this mutual wronging case compared to a non-mutual analogue, where similar wrongs are committed against third parties. Here is a non-mutual analogue:

Battery:

Blake attacks Ciara just for fun, injuring her. Amy attacks Djamel just for fun and injures Djamel's arm as badly as Blake injured Ciara's arm. Blake blames Amy for her wrongful attack. Other things are equal.

It is not as though when Blake blames Amy in this case, he blames her *more* hypocritically than he does in the first case. Hypocrisy, as I argued, is a function of similarity in culpable wrongdoing, plus the extent to which the blamer appropriately acknowledges wrongdoing.

These determinants of hypocrisy do not depend on who the victims of wrongdoing happen to be. So, the fact that Amy and Blake are victims of one another's wrongs does not seem to alter the degree of hypocrisy involved. In this respect, Co-Battery looks like just another instance of the same fault that we've been thinking about throughout.

A second corroborating observation is a scalar one. Is Blake's degree of standing to blame Amy any different between Battery and Co-Battery? Intuitively, it is not. To me at least, standing seems diminished to roughly the same degree. This leaves us with two instances of blaming which involve both a similar degree of hypocrisy, and a similar degree of standing. That may seem to further suggest that hypocrisy is what explains a loss in standing in mutual cases, just as it does so in non-mutual cases.

ii. Hypocritical non-victims

However, this is too quick. There has to be more going on, as is revealed by focussing on the observation that the blamer is a victim of wrongdoing in mutual wronging cases.

Notice that Blake is a victim of Amy's attack in Co-Battery, whereas he is not a victim in Battery; yet, we saw that his degree of standing is intuitively *similar* in both cases. This is puzzling. As mentioned, victimhood is a plausible instance of a *positive* standing-relevant fact; a fact that, on a scalar view, increases the procedural value of holding a wrongdoer accountable, and thus increases the moral appropriateness of blaming. But, if this is right – if victimhood does increase standing – we would normally expect Blake's standing in the mutual case to be increased relative to the non-mutual case. Although he is equally hypocritical in both cases, and this diminishes his standing, he is a victim in the mutual case, and not a victim in the other. So, why does this factor not give him greater standing?

One possibility is something I explored in the last chapter: that Blake's hypocrisy cancels the impact of his victimhood in the mutual case. Though his victimhood would ordinarily enhance his standing, his hypocrisy cancels out that positive impact. But that, I also pointed out, is not plausible. The significance of victimhood does not depend on whether we have committed similar wrongs, or whether we acknowledge those wrongs. The reason I gave

for this is that the value of a victim holding a wrongdoer accountable is grounded in the ability of the wrongdoer to deliver the responses she owes to the very person to whom she owes them. The fact that that person herself commits similar wrongs or doesn't acknowledge this does not affect the wrongdoer's corrective duties, and thus does not undermine the value of her engaging in these interactive responses.

The point that hypocrisy does not in itself disable victimhood comes out more clearly when we compare Battery with another variant, Battery 2:

Battery:

Blake attacks Ciara just for fun, injuring her. Amy attacks Djamel just for fun and injures Djamel's arm as badly as Blake injured Ciara's arm. Blake blames Amy for her wrongful attack. Other things are equal.

Battery 2:

Blake attacks Ciara just for fun, injuring her. Amy attacks Blake just for fun and injures Blake's arm as badly as Blake injured Ciara's arm. Blake blames Amy for her wrongful attack. Other things are equal.

Blake blames equally hypocritically in both cases; but, intuitively, he has more standing in the second. Since victimhood has a positive impact in the second, equally hypocritical case, hypocrisy does not disable the importance of victimhood by itself.

Since hypocrisy does not disable victimhood, we need a more complicated explanation of why Blake lacks standing in *Co-Battery*. Some further fact is in play in mutual wronging cases, which either cancels, or, in addition to hypocrisy, countervails, the fact that Blake is a victim. What is this further fact?

2. Impairment

Here is a way of putting the problem which hints at an answer. We need to understand why, despite being a victim, my standing to blame you for wronging me when I similarly wronged you is no higher than my standing to blame you when I am not a victim, and have similarly wronged someone else. Perhaps the distinguishing factor, then, is just this: that *I* similarly

wronged *you*. We can describe the fact that Blake similarly wrongs Amy as another *connective fact* – a fact that connects the blamer and the blamed conduct in a certain way. What is it that connects Blake to Amy and the wrongful attack? Well, it is not only the fact that he is the victim of her attack, and that his own attack is morally similar; it is also the fact that *Amy* is the victim of *his* attack. This indicates that victimhood and wronging can have a converse effect on standing: whereas being a victim of a person’s wrongdoing can be a standing-enhancer, wronging a person – making that person a victim of our wrongdoing – can be a standing-diminisher. A plausible explanation for this seems *relational*. Wronging a person, it seems, has an *impairing* effect on our relationship with that person, in a way that affects our standing to hold them accountable. However, we should be careful about how we cash out the idea of relational impairment.

i. The Difference View

It is commonplace to understand relational impairment as an upshot of wronging: wronging, some argue, changes a moral relationship between two agents, in a way that impairs this relationship.² But this proves to be a somewhat elusive thought. To understand why we should think of wronging as implying impairment, we need to be more precise about what a ‘moral relationship’ ultimately is. However, it is difficult to articulate what a moral relationship is in an informative way.

A very minimal way of understanding a moral relationship is just as an assessment of the difference between *directed duties* and behaviour: between, on the one hand, what agents owe to each other, and, on the other, how they behave. For example, we might call a moral relationship ‘good’ when the difference between how we behave and what we owe to someone

² Indeed, for some, impairing a moral relationship is *what it is* to wrong a person (See e.g. Frick, forthcoming). Also compare Driver (2017), who thinks about wronging as violating ‘relationship ideals’ (pp. 212-213). We might also think of my view below as a view about relationship ideals, though the difference is that the ideals on my view, can govern a future relationship that we might enable, whereas the ideals Driver has in mind govern a pre-existing relationship.

is morally *positive*, such as when I improve your well-being without owing that improvement to you, and without doing anything I owe it to you not to do, such as violating your rights. Similarly, we might call a moral relationship ‘okay’ if there is *no* difference between what parties owe and how they behave, such as when I don’t improve your wellbeing in a way that I don’t owe it to you to do, or when I do improve your wellbeing in a way that I owe it to you to do, without violating your rights in the process. Last, we might call a moral relationship ‘not okay’ when the difference between how parties behave and what they owe is morally *negative*, such as when I do something I owe it to you not to do (like attacking you for fun), or when I don’t do something I owe it to you to do (like failing to benefit you if I easily can). These negative differences are the result of parties wronging each other. So, we might define an impaired moral relationship as its having a ‘not okay’ status. Let us call this the *Difference* view.

Even if true, the Difference view cannot be a complete account of the moral relationship that is meant to be impaired by wrongdoing. To claim that an impaired relationship involves a change in our assessment of differences between directed duties and behaviour only assumes that there *is* such a relationship in the first place; one which predates the wrongdoing, and has been impaired by the wrongdoing. Assuming that relationship, though, does not tell us what it is.

Furthermore, it’s not clear why we should think about the impaired thing as a *relationship* at all.³ Ordinarily, by relationships, we mean something *historical*: we typically assess the status of people’s relationships based on how they interact with each other over a certain period of time, and on certain contingent facts about both of them, such as whether they agree, or promise, to do things for each other, or whether they are both involved in a practice or institution which involves certain roles, like a teacher and student, or manager and employee.

³ We might distinguish a relationship from a *relation* – perhaps there is a moral relation between people who are strangers, which is something thinner than a moral relationship. I myself am not sure how to explain the difference between ‘relational impairment’ and ‘relationship impairment’ in a way that would avoid the sort of concern I raise.

Wrongs against a person can be created by all of these things: a teacher can wrong a student by being a bad teacher, and I can wrong someone by not doing what I promised. However, wronging a person *in general* does not require any of these contingent facts. I can wrong someone with no particular historical connection to them at all – for example, I can wrong you, a stranger, by stealing your things.⁴ But it is odd to suggest that my doing this has any effect on my *relationship* with you. There was no relationship to begin with.

No doubt, we may be able to enjoy certain bonds with others without much historical connection with them. One bond that I have argued that hypocrisy forestalls, for example, is a solidaristic bond between ‘co-agents’ who are similarly morally situated: being co-agents does not require a particular history of interaction; it only requires shared capacities, and similar wrongdoing. However, we should distinguish between things that *create the conditions* for a certain relationship, on one hand, and things that are needed to *instantiate* that relationship on the other. What creates the conditions for a bond of solidarity – i.e. being a co-agent – does not require historical interaction. What instantiates the bond, though, is public mutual recognition of this fact. Public mutual recognition does require an interaction. But as the theft example above shows, wronging people doesn’t generally require interacting with them.

Perhaps attitudes are sufficient for a relationship. Another type of relationship discussed in connection with hypocrisy is a relationship of moral equality: a possible view is that my attitudes towards people can impair my relationship of moral equality with them. For example, I can wrong members of a certain group by holding attitudes of superiority towards that group without ever interacting with any of these members; perhaps holding these attitudes by itself impairs my relationship of moral equality with all those people. But, still, this would not explain the basic moral relationship that wronging *in general* is supposed to impair. In general, it is possible to wrong people without having any attitudes towards them – we might owe benefits to people far away in the world, whom we’ve never thought about at all, and fail to give them what we owe. Since wronging in general, unlike the wrong of unequal regard,

⁴ Jay Wallace calls this class of wrongdoing ‘monadic’ (2010, p, 340).

need not involve particular attitudes, it is not clear why a relationship that wronging impairs would involve any attitudes either. So, again, what does it involve, and why call this thing a relationship? Indeed, what does introducing the idea of a moral relationship add beyond restating the fact that agents can wrong each other?⁵

ii. *Attitudes we owe*

Scanlon (2008) has offered a way to understand moral relationships that does not require people to have attitudes toward one another.⁶ However, it is hard to distinguish Scanlon's view from the Difference view.

Scanlon agrees that our relationships with others are affected by the attitudes we might take toward one another. But he argues that this does not require that people actually hold certain attitudes. Relationships, Scanlon thinks, are constituted not only by the attitudes people actually have, but also by the attitudes they ought to have. He writes:

[M]orality requires that we hold certain attitudes towards one another simply in virtue of the fact that we stand in the relation of "fellow rational beings." [...] These attitudes and dispositions define what I am calling the moral relationship: the kind of mutual concern that ideally, we all have towards other rational beings (2008, p. 140).

This thought implies that there is a genuine relationship that I impair with someone I wrong even if I never think about that person. Even if I am totally unaware of and never consider certain people, it is plausible that I *ought* to have shown greater concern for these people in the way I behaved: I ought to have bothered to think about the implications of my actions, and to have been aware that these particular people exist, and have claims on me. I lacked the right attitudes towards these people. Since Scanlon holds that our moral relationships with others are a matter of the attitudes we ought to have, lacking attitudes we ought to have is sufficient, on his view, to impair a moral relationship.

⁵ Thanks to Adam Slavny for discussion on this point.

⁶ Scanlon (2008, ch. 4)

But this reduces to a version of the Difference view. Our challenge, remember, is to describe *what it is* that's impaired by a difference in how we behave and what we owe to others. Scanlon's thought points to a further difference of this kind, only one involving attitudes rather than behaviour: viz., a difference between the attitudes we have and the attitudes we owe it to others to have. But, again, saying we impair a moral relationship by not having the attitudes we owe each other is just another way of saying that we impair a relationship by wronging each other. This tells us nothing about the nature of the relationship that is impaired by defective attitudes.

Rahul Kumar (2015) suggests a response:

[The] moral relationship holds between any two individuals in virtue of a vulnerability to one another's attitudes, one shared by those who are both capable of, and care about, leading rationally self-governed lives [...] This mutual vulnerability is the grounds for holding all individuals to stand in a certain kind of relationship to one another, the moral relationship (2015, pp. 258-259).

Once again, though, this more or less re-describes Scanlon's point. It is true that regardless of any attitudes or historical connections we have, we as agents are mutually vulnerable to one another's attitudes. That is implicit in what was just noted above: that we might wrong one another by lacking certain attitudes towards them. The actual substance of this relationship of mutual vulnerability – beyond getting us to see the possibilities of wronging each other – remains unclear.

This is not, of course, to imply that there is no such substance; only that it is very difficult to pinpoint. In the light of this difficulty, I propose a different way to explain the significance of wronging to relationships between strangers.

iii. Prospective relationships

My suggestion is closely related to Scanlon's, but it does not assume a pre-existing relationship. Instead, I claim that wrongdoing can impair any *prospective* relationship that a victim and wrongdoer might form.

Like Scanlon and others, I think that wronging a person involves a deficiency in the value of our relationship with that person.⁷ However, this deficiency need not be a deficiency in a relationship which pre-dates the wronging. Rather the deficiency might be in a possible relationship that we may now form with the victim: uncorrected, any relationship I might now have with someone I wrong is not as valuable as it would otherwise be. We can see this in examples where people form relationships post-wronging; indeed where wronging is the inception of a relationship. Consider the movie *Parasite*, where each member of one family feigns false identities to convince members of another family to hire them for various domestic roles. It happens that both families are delighted by the arrangement. Despite this, the audience is uneasy about the situation, and I think part of that unease lies in the knowledge that the relevant relationships will never realise their ideal form unless the initial deception is corrected. Although there were no relationships between the various family-members prior to the wrongs, then, any relationship formed thereafter is impaired in virtue of them.

The upshots of this alternative view of impairment may be similar to Scanlon's view of blame (in Chapter 2, I called it *distancing blame*), and blame's warrant. Scanlon argues that if I know someone has wrongly stolen from me, I am warranted in judging that my relationship with this person is impaired. That judgment might manifest itself in further blaming attitudes and dispositions, such as intentions or desires not to interact with the person, or not to form agreements with her, or not to spend time with her. All these blame-constituting responses, Scanlon points out, invoke the idea of a relationship being impaired in the light of how the wrongdoer behaved. I agree on this much. But, notice the following difference between what I'm saying and what Scanlon says: to be warranted in blaming a person, I need not judge that a *pre-existing* relationship is impaired. Another thought that also justifies a conclusion about impairment is that, as things stand, in the light of what someone has done, any *future* relationship we might have with this person will be deficient in value.

⁷ Also compare Radzik (2014; 2009) and Duff (2010), whose account of standing and relationships I discuss below.

Scanlon's my own accounts also converge on the duties created by wronging someone: wronging comes hand in hand with duties to apologise, compensate, and so on, and if the person does not do this, we are warranted in judging our relationship to be impaired. There is clearly something relational about many of these demands. They demand that the wrongdoer interpersonally engage the victim to show recognition of the wrong, and apologise sincerely in the light of that recognition. The ideal form of this I have described as Radzik's 'victim-sees-wrongdoer-seeing-victim' dynamic; without that relational dynamic, full correction is lacking.⁸ Like Scanlon and Radzik, we could call the relevant demand *reparative*. But note, again, the difference: Scanlon thinks that these consequences of wrongdoing reveal the impairment of something pre-existing, which the wronging 'broke', as it were, and which the wrongdoer now needs to 'fix'.⁹ I do not rely on that thought. The minimal thing in a reparative demand is to make it the case that prospective interactions can proceed on ideal terms. Rectifying this deficiency *can* mean repairing a previously-broken relationship. But that is not all it can mean: it can also mean attempting to make possible the right kind of prospective relationship with someone we wronged, through the act of interpersonally addressing our wrongdoing with that person.

So much for my forward-looking take on relational impairment. Let us return to the issue of how this might affect the standing of a mutual wrongdoer to hold her victim accountable.

3. The Reparative Project

Recall our original example:

Co-Battery: Blake, a total stranger to Amy, attacks her just for fun, injuring her. Later, Amy's arm has recovered. Amy attacks Blake just for fun. The fun has nothing to do with exacting

⁸ A relational dynamic is not a relationship. But it is a necessary for a morally healthy relationship moving forward.

⁹ Radzik argues this (2014, p. 238)

revenge; at the time of the attack, Amy had forgotten about Blake's attack. She injures Blake's arm as badly as he injured Amy's arm. Blake blames Amy for her wrongful attack. Other things are equal.

If what I've said above is right, then when Blake wrongly attacks Amy, he impairs a prospective morally healthy relationship. Moreover, it is incumbent on Blake to create the conditions for a healthy relationship with Amy because this is part of what he owes her. I will now argue that his failure to do this is what diminishes his standing to blame.

i. Acknowledgment and repair

Blaming Amy for wronging him presents Blake with an appropriate opportunity to initiate a reparative relational project, one which he owes it to Amy to execute. But, where he blames her hypocritically – where he fails to acknowledge his wrongful attack on her – he fails to seize that opportunity. This explains his loss of standing.

Blake's failure to acknowledge in this case, notice, is not quite the same as it is non-mutual cases like Supremacists, or Battery. In those cases, I argued in Chapter 3, the failure to acknowledge similar wrongdoing forestalls a certain set of values that are advanced by mutual discussion about that type of wrongdoing. Now, that is *also* true in mutual wronging cases: since Blake does not publicly recognise his similar moral situation, the pair of them from cannot access the insights, solidarity and collaboration that mutually discussing their similar situations would achieve. However, in the mutual wronging case, the failure to acknowledge implies a further failure – it implies that Blake is also not delivering what he independently owes to Amy. The relationship Blake ought to seek to make possible, then, is more basic than a bond between similar wrongdoers, and one to which all agents (similar wrongdoers or not) are entitled to in our interactions: namely, one in which wrongdoing is corrected for. Making possible that basic relationship requires him to acknowledge his wrongdoing when he holds Amy accountable.

To be sure, Blake owes Amy more than merely interactive acknowledgment. As we've seen, his corrective duties also include e.g. compensation for Amy's injury, and ensuring that

he reforms his ways, regardless of whether he interacts with Amy. But, interactive acknowledgment is a core part of what he owes. We can see this by noting that Blake owes it to Amy to interactively *apologise* for his wrongdoing: not merely to internally realise what he did or desire to say sorry, but to communicate these attitudes in an eye-to-eye, Radzikian format. Blake cannot discharge that duty without acknowledging his wrongdoing – to apologise properly, we must fully and openly recognise the subject of our wrongdoing through an interaction with a person we wrong.¹⁰

What does this relational implication of Blake's failure to acknowledge have to do with his standing to blame? The Procedural Value account provides an answer. In Chapter 5, we saw that a core procedural value in blaming is to give wrongdoers a platform on which to understand, discuss and respond to their wrongs for the right reasons. In Chapter 6, and briefly above, I claimed that a similar value also explains why victims have more standing to blame wrongdoers than non-victims – because a victim's blaming is an appropriate moment for the wrongdoer to show understanding of, and deliver appropriate responses to, acting wrongly, directly to the person to whom these responses are owed, partly to ensure that future relationships between victims and wrongdoers can proceed in their ideal form. We can now see why this *same* set of procedural values – we might call them *reparative* values – explains why mutual wrongdoers have defective standing to blame their victims.

As suggested, the duty interactively to play our part in enabling an unimpaired prospective relationship is something Blake *independently* owes Amy. Irrespective of whether he blames her for her wrongdoing, or indeed whether she blames him, Blake owes Amy correction for his wrongdoing. Why does it matter, then, that he initiates a procedure of accountability? Again, a response draws on the Procedural Value account. Because a wrongdoer interactively owes the victim acknowledgment of wrongdoing, an accountability process, given its interactive nature, provides exactly the right kind of platform for discharging that duty. So, even though this duty is one that Blake already owes, it is particularly

¹⁰ Other than Radzik, also see e.g. Gardner (2018), Bovens (2009) and Bennett (2008; forthcoming).

appropriately discharged through the accountability process he initiates with Amy. In blaming Amy, therefore, Blake creates a particularly appropriate reparative opportunity; and yet, the way he interacts with Amy – by blaming her hypocritically, and thus not acknowledging wrongdoing – means that this chance of appropriate repair is lost. That loss in appropriateness is another way of describing his loss in standing.

We might wonder, though, whether this argument overlooks the *mutual* nature of the wrongs in such cases. We can't forget that the reparative duties above apply to *Amy* just as much as they do to Blake. Amy, after all, similarly wronged Blake. Thus, Amy's duties to engage in a reparative project are as strong as Blake's. This means that *her* failing to acknowledge her wrongdoing is equally anathema to the possibility of a future unimpaired relationship between the pair. Like mutual deliberation, a healthy relationship, whether current or prospective, is a two-way street. Blake acknowledging wronging Amy, by itself, achieves just a one-way street: only if Amy also acknowledges wronging Blake do we have a two-way street. As creating the conditions for an ideal prospective relationship is not Blake's sole responsibility, it might seem a mistake to attribute to him a failure to instantiate that procedural value.

However, as I have argued at various points, two wrongs don't make a right. Blake's reparative duty is not conditional on Amy's reciprocity; it is simply a consequence of what Amy is owed as an agent. Amy wronging Blake, even combined with Amy's reparative failure, does not affect her claims on Blake. So there is a procedural reparative value in Blake correcting for what he did regardless of whether Amy does the same, and if he blames her hypocritically, he is responsible for that loss in value.

Furthermore, the fact that Amy owes Blake something similar might only *strengthen* his reason to acknowledge; for it can help him get what he's owed.¹¹ The reason is that by acknowledging wrongdoing, Blake can model for Amy the response that she also owes to him. This echoes some of the motivational benefits of non-hypocrisy discussed in Chapter 3:

¹¹ Edwards (2018) has a similar thought (2018, n. 57)..

acknowledging our own wrongs sets the right example for those we blame, and can make them feel more comfortable about doing the same.

It might seem a bit strange to imagine Blake acknowledging wrongdoing partly to encourage Amy to reciprocate. One worry is that this undercuts the sincerity of Blake's acknowledgment; another is that Amy's apology is insincere if she is coaxed into it by Blake. Perhaps we would be uneasy about Blake having the goal of Amy's reciprocity in mind. But, that need not be goal. We can sincerely acknowledge our wrongs and hope that, as a side-effect, others will also acknowledge their wrongs because they have witnessed a good example of how to do this, and why doing it is important, rather than because they are cajoled or pressured into correction. Blake has a special reason to provide such an example, given that he is Amy's victim. Hence another reason for him to acknowledge.

What mutual wronging cases show, then, is that although the *degree* of hypocrisy may be the same in different cases of blame (where the degree depends on similarity of wrongdoing and culpability, plus the authenticity of acknowledgment), the *fault* of the same degree of hypocrisy can differ in gravity. Sometimes acknowledgment can achieve procedural value over and above the values of mutual discussion between similar wrongdoers. Where that is true, the failure to acknowledge comes at a greater detriment to the blamer's standing than the same failure in other instances of similar wrongdoing. And mutual wronging cases are a prime example. Compared to mutual discussion between similar wrongdoers, the value of mutual discussion between similar *mutual* wrongdoers is greater, because it also allows two parties, both of whom are owed certain responses and owe certain responses, by and to each other, to try to deliver those responses, and ensure that future relationships could proceed unimpaired.

We now have a way to make sense of our original problem: the problem of explaining why, in Co-Battery, Blake's standing is similar in its degree to Battery, despite his being a victim of wrongdoing in the mutual case, and despite his being equally hypocritical in both cases. The difference in Co-Battery is that although Blake is a victim, he is also blaming his victim. His failure to acknowledge wrongdoing thus has a stronger negative effect on his standing than the same failure in non-mutual cases, because that failure is a greater procedural loss. To put

it another way: hypocrisy in mutual wrongdoing cases more strongly countervails the positive impact of victimhood than it does in non-mutual cases.

ii. Historical views

This forward-looking account of standing in mutual wrongdoing cases – one that focusses on the value of relationships that accountability processes might now build, rather than on what happens to the value of pre-existing relationships – gains further credibility when we consider problems with alternative views, which have the opposite focus. I have already pointed out that the idea of a general pre-existing moral relationship is difficult to explain without merely restating the possibility of wrongdoing one another. But the idea that such a pre-existing relationship grounds *standing* in these cases is additionally problematic. I will first explain these problems, and then offer a diagnosis.

a. Problems

One way in which impairment of a pre-existing moral relationship is taken to affect standing is by affecting the warrant of the judgments that blame involves. This is how Scanlon explains standing in the mutual cases (2008, pp. 175-206). As mentioned above, for Scanlon, a central judgment in blaming is that the blamee, by behaving in a certain way, impaired her moral relationship with the blamer. To repeat, this judgment may not be a consciously formed: Blake's blame of Amy for attacking him might manifest in other kinds of revisions in his attitudes and dispositions, such as intentions to defend himself around Amy, or dispositions no longer to trust her, or desires not to spend time with her at all. But ultimately, what explains Blake's revised attitudes and dispositions is a cognitive judgment – a judgment that Blake's reasons to spend time with, or trust, Amy are not as powerful as before; and this is tantamount to judging that the relationship is impaired. In the mutual wrongdoing case, however, Scanlon argues that this judgment – that Amy impaired his relationship with her – is mistaken. He argues this for the following reason. Remember that in Co-Battery, Blake attacked Amy *before* she attacked him. Thus, it was *Blake*, and not Amy, who made it the case that either of them

had any reason to distance themselves from the other. This means it was Blake, and not Amy, who impaired the relationship; for him to contend that Amy did so is a case of mistaken identity. That, Scanlon concludes, is why Blake lacks standing: because his blame involves a mistaken judgment about who impaired the relationship, he lacks standing to blame.

There are several problems with this argument,¹² but let me focus here on a core issue. Briefly imagine some alternative chronologies of wrongdoing. Consider Blake's standing to hold Amy accountable when the parties wrongly attack each other at the same time (*Brawl*). Or consider Blake's standing to hold Amy accountable as a second impairer: Amy attacks him first, and he attacks her later (*Co-Battery 2*). In all these cases, his deficiency in standing is intuitively as low as it is in *Co-Battery (1)*, and low for the same kind of reason. But in neither case is he the original impairer; so in neither case does he commit the mistaken identity Scanlon has in mind. This suggests the following. The set of historical facts that Scanlon focusses on – facts about the chronology of impairment – are altogether immaterial to our assessment of standing. It's true in *Co-Battery (1)* that Blake was the first impairer, but that is in fact just a contingent feature of the case. Given that standing seems equally poor when the chronology of impairment changes, that contingent historical feature fails to explain why Blake's standing in the original case is diminished.

¹² One problem is that it is an argument about the wrong issue (i.e. about *warrant*, rather than about *standing*) – this is my interpretation of Kasper Lippert-Rasmussen's and Jay Wallace's concerns about Scanlon's view (Lippert-Rasmussen, 2013, p. 310; Wallace 2010, pp. 320-321). Another is that the argument assumes that blame implies a judgment that the blamee is the *first* impairer of a relationship, which looks like an arbitrary stipulation. It seems equally fine and potentially blame-constituting to judge e.g. that blame involves the judgment that the blamee is *an* impairer of a moral relationship, given that relationships, plausibly, can be repeatedly impaired by different parties. If this is right, the core judgment in Blake's blame is probably true, rather than false: since Amy wrongly attacks Blake, she is *an* impairer of her and Blake's relationship. Thus, the warrant of Blake's judgment does not explain a loss of standing.

In contrast, the account I gave – that standing is diminished by a failure to create the conditions of an unimpaired prospective relationship – explains the relevant deficiency in all three cases, and in an appropriately uniform way. Whether Blake wrongs Amy first, second, or simultaneously, the demand on him to make possible an unimpaired relationship with her, through acknowledging wrongdoing, is equally stringent either way. So, his failure to meet that demand is a loss in procedural value either way.

A second alternative account of the relationship between impairment and standing holds that standing is directly grounded in these relationships, and, therefore, that defective standing is directly explained by their impairment. Duff (2010) argues this (2010, pp. 138-140). Indeed, one way to understand Duff's view is as follows: standing to blame *just is* to stand in the right kind of relationship with a person that makes holding them accountable apt, and that this is no less true of relationships between moral agents as it is of, say, friendships or partnerships. Like Scanlon, Duff posits that while we all have more intimate relationships with some than others, all of us bear a moral relationship to one another in virtue of our agency; in other words, our agency places us in relationships as co-members of a moral community. This basic moral relationship is ultimately what grounds the standing of any agent to hold another accountable. However, these standing-grounding relationships are impaired, Duff argues, when agents wrong one another: in wronging a person, we fail to treat that person as a co-member of a moral community; we effectively 'exclude' them from that community in the light of how we treat them. This means that the basic standing-grounding relationship is lost: we no longer relate to those we wrong as co-agents. As that basic relationship is what's needed to make accountability apt in the first place, its impairment means that we no longer have standing to engage in these practices with those we wronged, at least unless we acknowledge wrongdoing.¹³

¹³ For a similar interpretation of Duff's argument, see Watson (2015). Also compare Mason (2019, ch. 4-6) for the idea that conversational accountability is restricted to co-members of a moral community, and Upadhyaya (2020) for criticism of that idea.

One problem with this account is of relevance to us here. The account struggles to explain the insight that gives rise to the Worse-than-Nothing Problem – namely, that accurate hypocritical blaming can be more morally apt than no accountability process at all. To see this, consider again the three options available to Blake prior to blaming Amy in Co-Battery:

- a. don't hold Amy accountable at all;
- b. hold Amy accountable hypocritically; and
- c. hold Amy accountable non-hypocritically (by acknowledging having wronged her).

Duff's account is able to explain why the third option is more morally apt than the second. We might for example think that in acknowledging his wrongdoing, Blake *restores* the moral relationship he previously impaired by wrongly attacking Amy. Perhaps that now places him back in the standing-grounding relationship for apt accountability.¹⁴ However, the account cannot explain why the second option – hypocritical blaming – is morally more apt than the first option – not blaming. Unless Blake acknowledges his wrongful attack, it remains the case, on Duff's view, that the standing-grounding relationship is lacking. Thus, Blake cannot stand in the right relationship to hold Amy accountable at all. This result, I have argued throughout, is not palatable.

Again, the forward-looking view, which sees standing as a function of reparative prospects in blaming, attractively explains the appropriateness hierarchy above. Remember: Blake is himself a victim. Even if this fact about him is countervailed by his wronging Amy, it has some positive force: it is reparatively valuable not only that Blake apologises to Amy, but also that Amy apologises to him, since she also wronged him. If Blake does not engage Amy at all, an appropriate opportunity for Amy to issue an apology to Blake is lost. Compared to no accountability, then, hypocritical accountability is reparatively better. It is more procedurally

¹⁴ Duff gives this argument (2010, p. 139). I doubt it is right, for a reason similar to the one I offered against the Ledger view in Chapter 5: the badness of wrongdoing is weightier than the goodness of acknowledging wrongdoing. Similarly, the damage to relationships through wrongdoing may sometimes be weightier than the restoration of relationships that might be achieved by acknowledging wrongdoing.

valuable, and apt, that Blake engage Amy hypocritically rather than doing nothing. However, this option is also not as reparatively valuable as the third, non-hypocritical option, for reasons already given: the most reparatively valuable process involves both Blake and Amy playing their part in enabling a morally healthy relationship moving forward. For that, Blake must acknowledge his wrongdoing as well as urging Amy to do the same. Therefore, it is more apt that he initiates accountability and acknowledges wrongdoing rather than initiate accountability without doing so.

b. Diagnosis

Overall, it seems a mistake to explain the loss of standing of wrongdoers, and mutual wrongdoers, by appeal to a pre-existing moral relationship, even if a plausible account can be given of such a relationship. But what is the source of that mistake? The source, I think, is the assumption that standing in the context of a relationship between moral agents who share no history, works just like standing in the context of other sorts of relationships discussed in this thesis, such as partnerships, friendships, and so on. It is plausible that the latter kind of relationships can suffer impairments which make standing defective. But, once we pay attention to why this is so, we can see that the same reasoning does not apply to the former context. Let me explain.

I argued in Chapter 6 that intimate and collegiate relationships may be impaired in a manner relevant to standing: I reported the intuition, for example, that as partnerships break down – where a partnership has long been declining, say, or has been subjected to multiple acts of infidelity – these people’s standing to hold one another accountable wanes alongside that historical trajectory. The reason why can be extracted from the account I gave in Chapter 5 of why partners have special standing in the first place: special accountability between partners for certain kinds of conduct is part of the value of those historically-contingent relationships – accountability constitutes a healthy partnership, friendship, and so on. But where the relevant relationships become either more distant or less healthy, accountability practices do not instantiate the same degree of constitutive value as they otherwise would.

That is why the procedural value of blaming, and hence standing, diminish in correspondence with these historical facts.

However, the same cannot be said about the value of holding one another accountable *simply as moral agents*. The value that inheres in a process where agents who wrong each other mutually account for their wrongs does not diminish in correspondence with particular interactions, attitudes, and connections that these agents historically share. The reason is simple: our duties to correct for wrongdoing, and thereby enable a future morally healthy relationship, through an accountability process, does not depend on what happened to our relationships in the past. This particular set of duties stands the test of time, so to speak, so complying with them instantiates the same procedural value whatever the past.¹⁵ That, ultimately, is why the chronology of impairment, and the idea that wronging impairs a prior standing-grounding relationship, have little relevance. Whatever the past, we have duties to do something about it now and in the future, and standing is sensitive to how well we discharge those duties.

Conclusion

Recall that a broader aim of this chapter was to illustrate the wider theoretical appeal of the overall account of hypocrisy and standing developed in the thesis, by focussing on mutual wronging cases, and trying to understand what sets them apart from other cases of hypocritical blaming. We can summarise the discussion by summarising that appeal.

¹⁵ Here is an example which illustrates the distinction a bit further. I might seriously wrong my friend in such a way that leads her not to talk to me at all for decades. If she holds me accountable decades later, that does not, I think, constitute the value of our friendship. Still, I owe it to this person, decades later, as an agent to apologise if I hold her accountable for something similar, as the duty to apologise stands the test of time, and not doing so remains a procedural fault; and any prospects of rebuilding requires this.

The appeal of Scalar is that it illuminates a problem of standing raised by reflection on mutual wronging cases of hypocrisy, by allowing us to compare a blamer's degree of standing in mutual and non-mutual cases. Once we observe that the hypocrite's standing is no greater in mutual wronging cases than it is non-mutual cases where he is not a victim, a problem arises in explaining what other fact, beyond mere hypocrisy, makes a mutual wrongdoer's blame deficient, and why. I have said that the right response to this problem draws on the idea of relational impairment. However, we ought to explain the significance of impairment without being unduly historical – that is, without relying on the problematic idea that standing depends on the impairment of a relationship which pre-dates wronging.

The appeal of the Procedural Value account is that it offers precisely that kind of explanation. Irrespective of our prior relationships, building good relationship prospects with those we have wronged, and have been wronged by, remains procedurally valuable. Wrongdoers have deficient standing because they fail to advance these reparative prospects, despite creating an apt platform – i.e. an accountability process – on which to do so. That further feature of the case explains why, compared to the standing of an ordinary hypocrite who is not a victim, the standing of a hypocritical mutual wrongdoer to blame his victim is similarly deficient: his hypocrisy implies a bigger procedural fault, and thus countervails his victimhood to a greater degree.

Conclusion

Before philosophical discussion about wrongful hypocrisy became rife, Chris Bennett (2002) wrote that:

As with sex, blaming is an activity that superficially appears to consist in one person doing something to another—and indeed abortive or perverted forms of blame can be like this. On a deeper understanding, though, they both involve something that can only properly be done together (2002, p. 153).¹

This neatly encapsulates the core of this project.

At the beginning, I sought to vindicate the judgment that hypocrisy is morally wrong, and distinctively morally wrong. There was something about hypocrisy in itself that seemed off. Explaining these observations proved difficult for the Deception and Walk the Talk views discussed in Chapter 1. We may have reasons to practice what we preach, and to be forthcoming about the sort of person we are, but we need to understand our special reasons to behave non-hypocritically. I suspected we would find an answer by exploring hypocrisy in critique: something about engaging others on normative matters, it seemed, enables a distinctive wrong of hypocrisy. What we have found is that the distinctive wrong arises from a distinctive kind of disvalue: the disvalue, in Bennett's words, of not engaging 'properly'. Doing it properly means doing it *together*, and doing it hypocritically prevents us from fully doing it together.

A future project is to work through the implications of this idea for different forms of normative interaction, such as advice, warnings, and even praise.² But among the most

¹ On this point about sex, also compare e.g. Gardner (2018), MacKinnon (2016) and Nagel (1979, ch. 4).

² Hypocrisy in, and standing to, praise remains a neglected issue – exceptions are Lippert-Rasmussen (2013; forthcoming b) and Mason (2019, pp. 109-111).

pressing instances, and the paradigm we've focussed on, is blaming. Blaming is a response to a certain kind of behaviour that requires serious reflection in a person: namely, on what to do about their wrongdoing. Other forms of normative interaction might be about wrongdoing, but that feature is central to blaming. And the value of interpersonal normative interaction is also part of what makes blaming apt, as I argued in Chapter 2. The structure of blame involves one agent responding to a certain failure in another, and among the values that this structure can achieve is the implicated agents responding to these failures for the right reasons. These internal, or aptness-grounding, values of blaming are discursive: the special contribution blaming offers to our moral lives is the initiation of a conversational process, where people listen and respond to one another's other's views, and offer their own views, on the subject of wrongdoing and what to do about it. That requires two people to do something together.

A key philosophical oversight, though, is that these *same* values are poorly realised in hypocritical blaming. Perhaps that oversight owes partly to the Worse-than-Nothing Problem. Hypocrisy, as sceptics about its wrongness have stressed, does allow for a version of blaming together. When I talk about your wrongdoing with you, we can talk about it together, as both of us can listen and respond to one another on that subject. Thus, hypocritical blaming can be both positively valuable and apt. However, the point in Chapter 3 was that this is not the *best* version of blaming. A further set of values – those of *Mutuality* – involve you and I talking about *both* of our similar wrongs together. Blaming together properly is not simply a matter of whether we can address *a* subject together; it's about the *range* of subjects we might together address: we should be able together to address my wrongdoing as well as yours. The hypocrite forestalls that ideal, by failing to discursively acknowledge his own wrongdoing. And this explains what is distinctive about hypocrisy over and above deception and irrationality. Not only can the failure discursively to acknowledge wrongdoing be deceptive, and involve a mistaken application of reasons. The core problem is to hamper mutual deliberation among similarly situated agents. Some of the values of this ideal are instrumental – such as the additional moral knowledge and motivation that mutual deliberation can promote – but some are also non-instrumental – mutual deliberation constitutes the good outcomes where we

collaboratively shape our futures, and it is good in itself that wrongdoers form a bond through mutually recognising their morally similar situations. Future work will develop these ideals of Mutuality in more depth: considering, for example, whether collaborated outcomes are merely impersonally valuable; what exactly explains the significance of a solidaristic bond between wrongdoers; and what duties might constitute that bond. But, these ideals are the ones that I think hypocrisy distinctively undercuts. The fact that hypocritical blaming is a non-ideal version of blaming together explains how, despite being apt, its aptness is diminished, and it is in that sense inapt.

This insight furnished us with an answer to the Worse-than-Nothing Problem. It became apparent that hypocrisy might be wrong, not because it is the worst option for a blamer to pursue, but because it is second-best. The reason it is wrongly second-best is that is that the best type of blaming can be done at low cost. That, I argued in Chapter 4, is typically the case; sometimes the costs of acknowledging wrongdoing can be high, but normally, acknowledgment benefits hypocritical blamers, both by lowering the reaction costs of blaming, and by enabling blamers to gain morally from mutual deliberation about wrongdoing. This offers a different way of explaining why hypocrisy is non-instrumentally wrong. It is non-instrumentally important that we can justify our conduct to others. Yet, hypocrites can't do this when they blame where the personal costs of non-hypocrisy are low. The good consequences of non-hypocritical blaming for blamer and blamee play a role in making this true, but the wrongmaking issue is that we cannot justifiably deny those goods to others.

But that, we saw, is not the whole story of the moral significance of hypocrisy. We also needed to figure why the complaint about hypocritical blaming is cashed-out in a certain way: it's not simply that the blamer blames wrongly – that is also true when we blame innocent people, and when we overblame people. When it comes to the hypocrite, the complaint is different – it is not about blame's accuracy or proportionality, but about the blamer's standing; and this is a complaint we observed in other instances, such as entrapment, estrangement, and complicity. The latter half of the thesis considered what hypocrisy and these other instances of defective blaming have in common. Like the accomplice, entrapper and stranger, the

hypocrite's blame is deficient in what, in Chapter 5, I called Procedural Value. However, these different facts each have a distinct impact on procedural value. The particular procedural value that is lost in hypocrisy – of mutual discussion about wrongdoing between similar wrongdoers – is not exactly the same as the values that other facts forestall. Indeed, we saw a whole host of disvalues, and values, which depend on various facts connecting a blamer to the blamee and to the blamed conduct: the disvalue of getting a person to act wrongly with the intention of holding them accountable; the disvalue of not acknowledging a deeper connection that is involved in responsibility for token-identical wrongdoing; the value of accountability practices in intimate relationships, and the value of victims holding wrongdoers to account. Alternative explanations of a hypocrite's deficient standing – for example, their bad character, bad past, or denial of equality – fail to capture the procedural implications of these various facts. Future work will investigate in more depth how best to spell-out what counts as a procedural value; but I think losses and gains in procedural value are the mark of standing.

This brings to light a further standard for 'proper' blaming: it can be more or less morally apt that certain people, rather than others, participate in accountability processes together. This, Chapter 6 claimed, is the central normative implication of losses in standing. The fact that blaming a person is less procedurally apt because of something about me does not necessarily make my blaming a person morally wrong, or unauthoritative, or unentitled; though it's true that the hypocrite blames wrongly, that is not due solely to a procedural defect. It is also due to the low cost for the hypocrite of rectifying this defect. More generally, standing considerations themselves do not make a binary difference. They make a scalar difference, to the degree of appropriateness of blame, and to whether we ought to persuade certain people to blame rather than others, or to blame in one way rather than another. This will have implications for how we should think about standing to hold accountable in political contexts, such as the standing of states to try certain criminals. Exactly what the implications are is another topic for future investigation; but the general take-home is that rather than opting for no accountability process, we should ideally aim for maximal appropriateness in assigning accountability roles within public institutions.

We finished by examining a test case for the overall theory – the case of what I called mutual wronging in Chapter 7. Such cases reveal how the procedural fault of hypocrisy can differ in its gravity. When hypocrites blame their own victims, their hypocrisy is a particularly grave failure, as it prevents wrongdoers from correctively addressing victims. Mutual wronging is an instance of blaming where it is especially pressing for two people – one another’s wrongdoers and victims – to participate in accountability, and move forward together through that process. That more serious failure to blame ‘properly’ counts more heavily against the standing of a mutual wrongdoer. Once again, related issues need to be explored in developing the scalar and procedural value framework. One set of issues is how and to what extent the gravity of wrongdoing affects procedural value; another is whether someone could unilaterally diminish the value of accountability with another person by repeatedly wronging this person. However, the above illustrates how I aim to apply the theory moving forward.

What, ultimately, is the upshot of all this for how we should deal with hypocrisy in our everyday lives? Well, the sceptics are right about this much: hypocrisy and other defects in critique of others don’t always warrant the outrage they tend to attract. Whether defective because of hypocrisy or other facts, accurate blaming can still have a positive moral valence, and we should not attempt to ignore or silence accurate criticism. The sceptic’s mistake, however, is to assume that hypocrisy is therefore neither wrong, nor standing-abrasive. The tendency on all sides – sceptics on one hand, and ‘common sense’ on the other – has been to think that if hypocrisy is wrong, or creates a problem of standing, then it must have a negative moral valence, or erect a moral bar to blaming. That overall picture needs a rethink. Rather than implying that an act has a negative moral valence – i.e. that it is worse than nothing – norms of hypocrisy and standing emphasise a deficiency in moral value. Best understood, the norms don’t say – ‘Don’t blame!’ They say – ‘If you blame, blame properly!’ ‘Ergo’, I say, ‘do it together!’

References

- Alexander, L, & Ferzan, K., 2009. With Morse, S. *Crime and Culpability: A Theory of Criminal Law*. Cambridge: Cambridge University Press.
- Arpaly, N., 2002. *Unprincipled Virtue: An Inquiry into Moral Agency*. Oxford: Oxford University Press.
- Beade, G. Who Can Blame Whom? Moral Standing to Blame and Punish Deprived Citizens. *Criminal Law & Philosophy* 13(2), pp. 271-281.
- Bell, M., 2012. Coates, A. J, and Tognazzini, N. A. ed. The Standing to Blame: A Critique. *Blame. Its Nature and Norms*. New York: Oxford University Press, ch. 14.
- Bennett, C., forthcoming. Apology as a Normative Power.
- 2018. The Alteration Thesis: Forgiveness as a Normative Power. *Philosophy & Public Affairs* 46(2), pp. 2017-230.
 - 2008. *The Apology Ritual: A philosophical Theory of Punishment* (Cambridge: Cambridge University Press, 2008)
 - 2002. The Varieties of Retributive Experience. *Philosophical Quarterly* 52(207), pp. 145-163.
- Bovens, L., 2009. Must I Be Forgiven? *Analysis* 69(2), pp. 227-233.
- Bratman, M. ,2009a. Intention, Practical Rationality and Self-Governance. *Ethics* 119(3), pp. 411-433.
- 2009b. Intention Rationality. *Philosophical Explorations: An International Journal for the Philosophy of Mind and Action* 12(3), pp. 227-241.

- Bridges, J., 2009. Rationality, Normativity and Transparency. *Mind* 118(470), pp. 353-367.
- Broome, J., 2005. Have We Reason to Do as Rationality Requires? – A Comment on Raz. *Journal of Ethics and Social Philosophy* 1(1), pp.1-8.
- Brownlee, K., 2020. *Being Sure of Each Other: An Essay on Social Rights and Freedoms*. Oxford: Oxford University Press.
- 2019. Acting Defensively for the Sake of Our Attacker. *Journal of Moral Philosophy* 16(2), pp. 105-130.
 - 2016. The Lonely Heart Breaks: On the Right to Be a Social Contributor.
 - 2015. *What's Virtuous about the Law?* *Legal Theory* 21(1), pp. 1-21. *Aristotelean Society Supplementary Volume* 90(10), pp. 27-48.
- Chang, R., 2020. Do We Have Normative Powers? *Aristotelean Society Supplementary Volume* 94(1), pp.. 275-300.
- Cohen, G. A., 2013. Otsuka, M. ed. *Finding Oneself in the Other*. Oxford: Princeton University Press, ch.6-7.
- 1992. Peterson, G. B., ed. Incentives, Inequality and Community. *The Tanner Lectures on Human Values, Volume Thirteen*. Salt Lake City: University of Utah Press, pp. 262-329.
- Cramwinckel, F. M., van Dijk, E., Scheepers, D., and van den Bos, K., 2013. The threat of moral refusers for one's self-concept and the protective function of physical cleansing. *J. Exp. Soc. Psychol.* 49(6), pp. 52–67.

Crisp, R. & Cowton, C., 1994. Hypocrisy and Moral Seriousness. *American Philosophical Quarterly* 31(4), pp. 343-349.

Dancy, J., 2004. *Ethics Without Principles*. New York: Oxford University Press.

Dakers, S., 2011. Why Ex-Offenders Should Be Given A Role In Cutting Youth Crime.

Guardian Online.

<<https://www.theguardian.com/society/joepublic/2011/sep/19/youth-justice-ex-prisoners-teaching>>

Darwall, S., 2006. *The Second-person Standpoint*. Cambridge, MA: Harvard University Press.

Dover, D., 2019. The Walk and the Talk. *Philosophical Review* 128(4), pp. 87-442.

Dolinko, D., 1991. Some Thoughts About Retributivism. *Ethics* 101(3), pp. 537-559.

Duff, A., 2018a. *The Realm of Criminal Law*. Oxford: Oxford University Press.

— 2018b. Responsibility and Reciprocity. *Ethical Theory & Moral Practice* 21, pp. 775-787.

— 2010. Blame, Moral Standing and the Legitimacy of the Criminal Trial. *Ratio* 23(2), pp.123-140.

Dworkin, G., 2000. Ullmann-Margalit, E. ed. Morally Speaking. *Reasoning Practically*. New York: Oxford University Press, pp.182-188.

Dworkin, R., 2011. *Justice for Hedgehogs*. Cambridge, MA: Harvard University Press.

Driver, J., 2017. Shoemaker, D. ed. Wronging, Blame and Forgiveness. *Oxford Studies in Agency and Responsibility Volume 4*. Oxford: Oxford University Press, ch. 10.

Ebels-Duggan, K. (2010). The Beginning of Community: Politics in the Face of Disagreement. *Philosophical Quarterly* 60(238), pp. 50-71.

Edwards, J., 2018. Standing to Hold Responsible. *Journal of Moral Philosophy* 16(4), pp. 437-462.
(Cited in text as pp. 1-26 from PDF download).

Enoch, D., 2014. Authority and Reason-Giving. *Philosophy and Phenomenological Research* 89(2), pp. 296 – 332.

— 2011. Giving Practical Reasons. *Philosopher's Imprint* 1(4), pp. 1 – 22

— 2009. Wouldn't It Be Nice If p , Therefore p (For a Moral p)? *Utilitas* 21(2), pp. 222-224.

Faulkner, P., 2007. What Is Wrong With Lying? *Philosophy & Phenomenological Research* 75(3), pp. 535-557.

Franklin, C. E., 2012. Coates, A. J, and Tognazzini, N. A. ed. Valuing Blame. *Blame: Its Nature and Norms*. New York: Oxford University Press, ch. 11.

Friedman, M., 2013. How To Blame People Responsibly. *The Journal of Value Inquiry* 47(3), pp. 271-284.

Frick, J., forthcoming. Dilemmas, Luck and the Two Faces of Morality.

— 2017. "[Zukünftige Personen und Schuld Ohne Opfer \(Future Persons and Victimless Wrongdoing\)](#)". Worauf es ankommt: Derek Parfits praktische

Philosophie in der Diskussion. Ed. Sebastian Muders, Markus Rüter, & Matthias Hoesch. Hamburg: Felix Meiner Verlag.

- 2016. What We Owe to Hypocrites: Contractualism and the Speaker-Relativity of Justification. *Philosophy & Public Affairs* 44(4), pp. 223-256.

Fricker, M., 2016. What's the Point of Blame? A Paradigm Based Explanation', *Noûs* 50(1), 165-183.

Fritz, K. G. & Miller, D., 2019. The Unique Badness of Hypocritical Blame. *Ergo* 6(19), pp. 545-569.

- 2018. Hypocrisy and the Standing to Blame. *Pacific Philosophical Quarterly* 97(2), pp. 118-139.

Gardner, J., 2018a. *From Personal Life to Private Law*. Oxford: Oxford University Press.

- 2018b. The Opposite of Rape. *Oxford Journal of Legal Studies* 38(1), pp. 48-70.

Harman, E., 2016. Morally Permissible Moral Mistakes, *Ethics* 126(2), pp. 366-393.

Herstein, O., 2020. Justifying Standing to Give Reasons: Hypocrisy, Minding One's Business and Knowing One's Place. *Philosophers' Imprint Early View*.

doi: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3490741

- 2017. Understanding Standing: Permission to Deflect Reasons. *Philosophical Studies* 174(12), pp. 109-132.

Hieronimi, 2004. The Force and Fairness of Blame. *Philosophical Perspectives* 18(1), pp. 115-148.

Holton, R., 2004. Rational Resolve. *Philosophical Review* 113(4), pp. 507-535.

- Horton, J., 2017. The All or Nothing Problem. *The Journal of Philosophy* 114(2), pp. 94–104.
- Howard, J., 2016. Moral Subversion and Structural Entrapment. *The Journal of Political Philosophy* 24(1), pp. 24–46.
- Hughes, P. & Warmke, B., 2017. Zalta, E. N. ed. Forgiveness. *Stanford Encyclopaedia of Philosophy*.
<https://plato.stanford.edu/archives/sum2017/entries/forgiveness/>
- Husak, D., 2011. Tonry, M. ed. Retributivism, Proportionality, and the Challenge of the Drug Court Movement. *Retributivism Has a Past: Has It a Future?* New York: Oxford University Press, ch. 12.
- Isserow, J. & Klein, C., 2017. Hypocrisy and Moral Authority. *Journal of Ethics & Social Philosophy* 12(1), pp. 191-222.
- Kagan, S., 1989. *The Limits of Morality*. Oxford: Clarendon Press.
- Kelly, E., 2018. *The Limits of Blame: Rethinking Punishment and Responsibility*. Cambridge, MA: Harvard University Press.
- King, M., 2020. King, M. Attending to Blame. *Philosophical Studies Early View*.
doi: <https://doi.org/10.1007/s11098-019-01260-w>
- 2019. Shoemaker, D. ed. *Oxford Studies in Agency and Responsibility Volume 6*. Oxford: Oxford University Press, ch. 11.
- 2015. Manipulation Arguments and the Standing to Blame. *Journal of Ethics & Social Philosophy* 9(1), pp. 1-20.
- Kittay, E. F., 1982. On Hypocrisy. *Metaphilosophy* 13(3-4), pp. 277-289.

Kolodny, N., 2008a. The Myth of Practical Consistency. *European Journal of Philosophy* 16(3), pp. 366-402.

— 2008b. Why Be Disposed to Be Coherent? *Ethics* 118(3), pp. 437-463.

— 2005. Why Be Rational? *Mind* 114(455), pp. 509-563.

Kumar, R., 2015. Clarke, R., McKenna, M. & Smith, A. M. ed. Contractualism and the Roots of Responsibility. *The Nature of Moral Responsibility: New Essays*. New York: Oxford University Press, ch. 11.

Lackey, J., 2013. Lies and Deception: An Unhappy Divorce. *Analysis* 73(2), pp. 236-248.

Lang, G, forthcoming. OVERRATING HYPOCRISY.

Leland, R. J., 2019. Civic Friendship, Public Reason. *Philosophy & Public Affairs* 47(1), pp. 72-103.

Lippert-Rasmussen, K., forthcoming a. Why the Moral Equality Account of Standing to Blame Fails. *Analysis*.

— Forthcoming b. On Praising Without Standing.

— 2018. *Relational Egalitarianism: Living as Equals*. Cambridge: Cambridge University Press.

— 2013. Kühler, M. & Jelinek, N. ed. Who Can I Blame? *Autonomy and the Self*, pp. 295-315.

Lister, Andrew, 2017. Public Reason and Reciprocity, *Journal of Political Philosophy* 25(2), pp. 155-172.

- MacKinnon, C. A. Rape Redefined. *Harvard Law & Policy Review* 10(2), pp. 431-77.
- Macnamara, C., 2015. Reactive Attitudes as Communicative Entities. *Philosophy & Phenomenological Research* 90(3), pp. 546-569.
- Maes, H., 2004. Modesty, Asymmetry, and Hypocrisy. *Journal of Value Inquiry* 38(4), pp. 485-497.
- Mason, E., 2019. *Ways to be Blameworthy: Rightness, Wrongness and Responsibility*. Oxford: Oxford University Press.
- Mason, M., 2003. Contempt as a Moral Attitude. *Ethics* 113(2), pp. 234-272.
- Matravers, M., 2006. Who's Still Standing? A Comment on Antony Duff's Preconditions of Criminal Liability. *Journal of Moral Philosophy* 3(3), pp. 320-330.
- McKenna, M. 2012a. *Conversation and Responsibility*, New York: Oxford University Press.
- 2012b. Coates, A. J, and Tognazzini, N. A. ed. Directed Blame and Conversation. *Blame: Its Nature and Norms*. New York: Oxford University Press, ch. 7.
- McKiernan, A., 2016. Standing Conditions and Blame. *Southwest Philosophy Review* 32(1), pp. 145-151.
- McKinnon, C., 2005. Hypocrisy, Cheating and Character-Possession. *Journal of Value Inquiry* 39(3-4), pp. 399-414.
- 1991. Hypocrisy, with a Note on Integrity. *American Philosophical Quarterly* 28(4), pp. 321-330.

McMahan, J. 2018. Woodruff, P. ed. Doing Good and Doing the Best. *The Ethics of Giving*, New York: Oxford University Press,

Merrick, R. (2019). 'Boris Johnson Cannot Criticise Us...' *The Independent Online*.
<https://www.independent.co.uk/news/uk/politics/boris-johnson-racism-football-hypocrisy-uefa-president-aleksander-ceferin-a9230391.html>

Murray, W. (2020). 'Wednesday Briefing: Trump Blames WHO and Freezes Funding.' *Guardian Online*.
<https://www.theguardian.com/world/2020/apr/15/wednesday-briefing-trump-blames-who-and-freezes-funding>

Minson, J. A., and Monin, B., 2011. Do-gooder derogation: disparaging morally motivated minorities to defuse anticipated reproach. *Soc. Psychol. Pers. Sci.*3(2), pp. 200–207.

Monin, B., Sawyer, P. J., and Marquez, M. J, 2008. The rejection of moral rebels: resenting those who do the right thing. *J. Pers. Soc. Psychol.* 95(1), pp. 76–93.

Moore, M., S., 1997. *Placing Blame: A Theory of the Criminal Law*. Oxford: Oxford University Press.

Nagel, T., 1998. Concealment and Exposure. *Philosophy & Public Affairs* 27(1), pp.3–30.
— 1979. *Mortal Questions*. Cambridge: Cambridge University Press.

Nelkin, D., 2016. Accountability and Desert. *Journal of Ethics* 20(1-3), pp. 173-189.

Nili, S., 2018. Integrity, Personal and Political. *Journal of Politics* 80(2), pp. 428-441.

- O’Neil, C., 2012. *Philosophy & Public Affairs* 40(4), pp. 301-333.
- Otsuka, M. & Voorhoeve, A., 2009. Why it Matters that Some Are Worse-off than Others: An Argument Against the Priority View. *Philosophy & Public Affairs* 37 (2), pp. 171-199.
- Otsuka, M. 2004. Skepticism about Saving the Greater Number. *Philosophy & Public Affairs* 32 (4), pp. 413-426.
- Owens, D., 2012. *Shaping the Normative Landscape*. Oxford: Oxford University Press.
- Pallikkathayil, J., 2019. *Philosophy & Phenomenological Research* 98(1), pp. 147-166.
- Parfit, D., 2017. McMahan, J. & Temkin, L. ed. Future People, the Non-Identity Problem, and Person-Affecting Principles. *Philosophy & Public Affairs* 45(2), pp. 118-157.
- 1982. Future Generations: Further Problems? *Philosophy & Public Affairs* 11(2), 113-172.
- PEA Soup, 2019a Neal Tognazzini. Discussion: Weaponizing Blame.
<http://peasoup.us/2019/10/weaponizing-blame-neal-tognazzini-for-ar-october/>
- 2019b. Discussion: The Puzzle of Hypocritical Self-Blame. David Shoemaker.
<http://peasoup.us/category/discussions/agency-and-responsibility/moral-responsibility/>
- 2017. Johan Frick. Discussion: Theron Pummer’s “Whether and Where to Give”.
<http://peasoup.us/2017/04/philosophy-public-affairs-discussion-pea-soup-theron-pummers-whether-give-critical-precis-johann-frick/>
- Pereboom, D., 2001. *Living Without Free Will*. Cambridge: Cambridge University Press.

- Pickard, H., 2013. Irrational blame. *Analysis*, 73(4), pp. 613-626.
- 2011. 'Responsibility without blame: empathy and the effective treatment of personality disorder'. *Philosophy, Psychiatry, Psychology* 18(3), pp. 209-224.
- Pummer, T., 2016. Whether and Where to Give. *Philosophy & Public Affairs* 44(1), pp. 77-95.
- Peter, F., 2017. Political Legitimacy. Zalta, E. N. ed. Political Legitimacy. *Stanford Encyclopaedia of Philosophy*.
<https://plato.stanford.edu/archives/sum2017/entries/legitimacy/>
- Radzik, L., 2014. Oberdiek, J., ed. Tort Processes and Relational Repair. *Philosophical Foundations of the Law of Torts*. Oxford: Oxford University Press, ch. 11.
- 2012, On the Virtue of Minding Our Own Business. *Journal of Value Inquiry* 46, pp. 173-182.
- 2011. On Minding Your Own Business: Differentiating Accountability Relations within the Moral Community. *Social Theory & Practice* 7(4), pp. 574-59
- 2009. *Making Amends: Atonement in Morality, Law and Politics*. Oxford: Oxford University Press.
- Raz, J., 2005a, The Myth of Instrumental Rationality. *Journal of Ethics & Social Philosophy*, 1(1): 1–28
- 2005b, Instrumental Rationality: A Reprise, *Journal of Ethics * Social Philosophy*, Symposium 1, p. 119.
- 1989, Liberating Duties. *Law & Philosophy* 8(1), pp. 3-21.
- 1986. *Morality to Freedom*. Oxford: Oxford University Press.

- Riedener, S., 2019. The Standing To Blame, or Why Moral Disapproval Is What It Is. *Dialectica* 73(1-2), pp. 183-210.
- Renzo, M., 2012. Associative Responsibilities and Political Obligation. *Philosophical Quarterly* 62(246), pp. 106–127.
- Roadevin, C., 2018. ‘Hypocritical Blame, Fairness, and Standing.’ *Metaphilosophy* 49(1-2), pp. 137-152.
- Rosen, G., 2015. Clarke, R., McKenna, M. & Smith, A. M. ed. The Alethic Conception of Responsibility. *The Nature of Moral Responsibility: New Essays*. New York: Oxford University Press, ch. 3.
- Runciman, D., 2008. *Political Hypocrisy*. Princeton: Princeton University Press.
- Scanlon, T.M., 2008. *Moral Dimensions: Permissibility, Meaning, Blame*. Cambridge, Mass: Belknap Press.
- (1998) *What We Owe to Each Other*. Cambridge, Mass: Harvard University Press.
- Scheffler, S., 2002. *Boundaries and Allegiances: Problems of Justice and Responsibility in Liberal Thought*. Oxford: Oxford University Press.
- Setiya, K., 2007. Cognitivism about Instrumental Reason. *Ethics* 117(4), pp. 649-673.
- Shabo, S., 2012. Where Love and Resentment Meet. *Philosophical Review* 121(1), pp. 95-124.
- Sher, G., 2006. *In Praise of Blame*. Oxford: Oxford University Press.

- Seim, M., 2019. The Standing to Blame and Meddling. *Teorema* 38(2), pp. 7-26.
- Shelby, T., 2016. *Dark Ghettos: Injustice, Dissent, and Reform*. Cambridge, MA: Harvard University Press.
- Shiffrin, S., 2014. *Speech Matters: On Lying, Morality and the Law*. Princeton: Princeton University Pres.
- Shoemaker, D., 2018. Cherry, M. & Flanagan, O. ed. You Oughta Know! Defending Angry Blame. *The Moral Psychology of Anger*. London: Rowman & Littlefield.
- Shoemaker, D., & Vargas, M., 2019. Moral Torch-Fishing: A Signalling Theory of Blame. *Nous Early View*, pp. 1-22. doi: <https://doi.org/10.1111/nous.12316>
- Sliwa, P., forthcoming. Reverse Engineering Blame. *Philosophical Perspectives*. Available at: http://paulinasliwa.weebly.com/uploads/1/9/0/4/19046427/reverse_engineering_blame_f_or_phil_perspectives.pdf
- Smilanksy, S., 2006. The Paradox of Moral Complaint. *Utilitas* 18(3), pp. 284-290.
- Smith, A. M., 2013. Coates, A. J, and Tognazzini, N. A. ed, Moral Blame and Moral Protest. *Blame. Its Nature and Norms*. Oxford: Oxford University Press, ch. 28-48.
- 2007. On Being Responsible and Holding Responsible. *The Journal of Ethics* 11(4), pp. 465-484.
- Smith, M., 1995. *The Moral Problem*. Oxford: Blackwell.

- Smith, P., 2007. Frey, R. G. & Wellman, C. H. ed. Bad Samaritans, Acts and Omissions. *A Companion to Applied Ethics*. Oxford: Blackwell, ch. 35.
- Szabados, B. & Soifer, E., 2004. *Hypocrisy: Ethical Investigations*. Peterborough, N. H.: Broad View Press.
- 1999. Hypocrisy, Change of Mind and Weakness of Will: How to do Philosophy with Examples. *Metaphilosophy* 30(1-2), pp. 60-78.
- Statman, D., 1997. Hypocrisy and Self-Deception. *Philosophical Psychology* 10(1), pp.73-79.
- Strawson, P.F, 1962. Freedom and Resentment *Proceedings of the British Academy* 48, pp. 187–211
- Swift, A., 2004. The Morality of School Choice. *Theory and Research in Education* 2(1), pp. 7-21.
- Tadros, V., 2020a. Criminalization: In and Out. *Criminal Law & Philosophy. Early View*.
doi: <https://doi.org/10.1007/s11572-020-09536-7>
- 2020b. Distributing Responsibility. *Philosophy & Public Affairs* 48(3), pp. 223-261.
- 2016. *Wrongs and Crimes*. Oxford: Oxford University Press.
- 2009. Poverty and Criminal Responsibility. *Journal of Value Enquiry* 43(3), pp. 391–413.
- Todd, P., 2017. A Unified Account of the Moral Standing to Blame. *Noûs* 53(2), pp. 347-374.
- 2012. “Manipulation and Moral Standing: An Argument for Incompatibilism,” *Philosophers’ Imprint* 12(7), pp. 1 – 18.
- Tognazzini, N. A., 2020. Silence and Silence: On Being Judgmental. Schmidt, S. & Gerhard E. ed. *The Ethics of Belief and Beyond: Understanding Mental Normativity*. New York: Routledge, ch. 13.

Tomlin, P., 2014. Retributivists! The Harm Principle Is Not For You! *Ethics* 124(2), pp. 272-298.

Tosi, J. and Warmke, B., 2020. *Grandstanding: The Use and Abuse of Moral Talk*. New York: Oxford University Press.

— 2016. Moral Grandstanding. *Philosophy & Public Affairs* 44(3), pp. 197-217.

Turner, D., 1990. Hypocrisy. *Metaphilosophy* 21(3), pp. 262-269.

Upadhyaya, forthcoming a. Hypocrisy: Good, But Wrong. (under review)

— Forthcoming b. Standing, on a Scale. (under review)

— 2020. Review of Elinor Mason, *Ways to Be Blameworthy*, *Ethics* 130(3), pp. 455-460.

Wallace, R. J., 2019. Coates, A. J, and Tognazzini, N. A. ed, Moral Address: What it Is, Why it Matters. *Oxford Studies in Agency and Responsibility Volume 5*. Oxford: Oxford University Press, ch. 4.

— 2011. Dispassionate Opprobrium. *Reasons and Recognition: Essays on the Philosophy of T. M. Scanlon*. New York: Oxford University Press, ch. 15.

— 2010. Hypocrisy, Moral Address and the Equal Standing of Persons. *Philosophy & Public Affairs* 38(4), pp. 307-341.

— 1994. *Responsibility and the Moral Sentiments*. Cambridge, MA: Harvard University Press.

Warmke, B., 2016. The Normative Significance of Forgiveness, *Australasian Journal of Philosophy* 94(4), pp. 687-703.

Watson, G., 2015. A Moral Predicament in the Criminal Law. *Inquiry* 58(2), 2015, pp. 168-188.

Wellman, C. H., 2017. *Rights Forfeiture and Punishment*. Oxford: Oxford University Press.

Wertheimer, R., 1998. Constraining Condemning. *Ethics* 108(3), pp. 489-501.

Wolf, S., 2011. Wallace, R. J., Kumar, R. and Freeman, S. ed. *Reasons and Recognition: Essays on the Philosophy of T. M. Scanlon*. New York: Oxford University Press, ch. 14.