

The Epistemology and Ethics of Epistemic Partiality in Friendship

Jack Robert Warman

PhD

University of York

Philosophy

May 2019

The Epistemology and Ethics of Epistemic Partiality in Friendship

Abstract

You should give your friend the benefit of the doubt if you hear nasty rumours about them. And you should look on the bright side when you think about their good qualities. At least, that's what popular wisdom teaches us about friendship. These practices are just two examples of the phenomenon of 'epistemic partiality'. We seek evidence that supports a favourable view of our friends, and we avoid evidence that challenges that view. We require more evidence to form unfavourable beliefs about them, and we require less to form favourable beliefs about them.

In this thesis, I argue that we should treat our friends with *limited* epistemic partiality. Unless the moral or prudential stakes are too high, or your evidence is incontrovertible, you should show your friend epistemic partiality. Plausibly, close friendship is a caring relationship. A person's beliefs can harm and benefit other people, I argue, and our friends are especially vulnerable. Epistemic partiality protects our friends from such harm. Therefore, I conclude, we should treat our friends with epistemic partiality.

But there are limits. Is epistemic partiality epistemically irrational, so that we shouldn't treat our friends with epistemic partiality? I argue that even if it's epistemically irrational to treat someone with epistemic partiality, that doesn't mean we shouldn't do so. Is epistemic partiality immoral? I'll consider to what extent moral considerations place limits on the epistemic partiality with which we should treat our friends. Finally, is the argument from care based on a mistaken conception of friendship? In responding to these objections, our understanding of friendship and epistemic partiality is further refined.

I conclude by reflecting on some possible lines of inquiry for the future. In particular, I suggest that it's possible we should treat the victims of sexual or racial discrimination or violence with epistemic partiality.

Table of Contents

Abstract	2
Contents	3
Acknowledgements	4
Author's Declaration	5
Introduction	6
Chapter One: Epistemic partiality in friendship	17
Chapter Two: Why we should treat our friends with limited epistemic partiality	37
Chapter Three: An epistemological objection to treating friends with epistemic partiality	72
Chapter Four: A moral objection to treating friends with epistemic partiality	101
Chapter Five: An objection to treating friends with epistemic partiality from the nature of friendship	125
Conclusion	141
References	147

Acknowledgements

First and foremost, I'd like to thank my family in the UK and in Chile for supporting me throughout my time at York and during my first year in Santiago. I'm especially indebted to Phil and Jill, my parents, to Kipp, my brother, and to Yeisil. They have always been there for me, even when they were thousands of miles away.

I'd also like to thank my friends for their encouragement and toleration. Special thanks are owed to Joshua Cockayne, Julie Kay, Christophe Robert, Hannah Goldwyn-Simpkins, Joe McGrail, Gordon Haynes, Andy Byrne, and John Packer. They've given me the benefit of the doubt more times than I care to remember.

In the Department of Philosophy at the University of York, I'm grateful to Christian Piller and Keith Allen for their feedback and comments on drafts of several chapters of this thesis. I'm extremely grateful to Janet Eldred, Carol Dixon, Karen Norris, and Jane Bertram for their support.

Finally, I must thank David Eford unreservedly. Without his encouragement and teaching, I never would have started my PhD. Without his guidance and insight, I never could have finished it. It is hard to exaggerate his contribution as a mentor and as a friend.

This work was supported by the Arts & Humanities Research Council (Award Reference: AH/L503848/1) through the White Rose College of the Arts & Humanities.

Author's Declaration

I, Jack Robert Warman, declare that this thesis is a presentation of original work and I am the sole author. This work has not previously been presented for an award at this, or any other, university. All sources are acknowledged as References.

Introduction

You hear a rumour about your friend. Usually it's easy to dismiss rumours as mere hearsay—nothing that would challenge your good opinion of your friend's character and conduct. This time, however, you can't help but worry that, just maybe, there's no smoke without fire. This uncertainty puts you in a very difficult position. If what you've heard about your friend is true, then you should think considerably less of them. You might even think twice about socialising with them so much in the future. But what if you're wrong? Well, if it turns out that you're misinformed but you form a harsh judgement about their character anyway, you could really harm them, and what's more, you might jeopardise a valuable friendship too. So, what should you do?

You should give your friend the benefit of the doubt in situations like this, at least according to popular wisdom. Actually, we have many ways to describe the sort of bias that we sometimes extend to our friends when presented with information that could diminish them in our judgement: 'She's your friend. Cut her some slack!' or 'Can't you just turn a blind eye?', we might say. At least this suggests that in day-to-day life, we're accustomed to the idea of treating our friends differently when we make judgements about them.

But what exactly does it mean to give someone the benefit of the doubt? Actually, giving the benefit of the doubt is one of many practices that have come to be known by concerned philosophers as 'epistemic partiality', owing to Sarah Stroud's (2006) influential paper 'Epistemic Partiality in Friendship'. (Along with Simon Keller's (2004) paper, 'Friendship and Belief', Stroud's paper can undoubtedly be credited with drawing the attention of analytic philosophers to this phenomenon.) Surely, epistemic partiality has something to do with marshalling your judgements and beliefs about them in a way that benefits them, or at least, in a way that doesn't harm them. Sometimes, treating your friend with epistemic partiality can involve withholding judgement about their character or conduct, especially in the light of new information. You might wait before you judge them, at least until you've been able to get their side of the story. In so doing, you seek to protect your good judgement of them. Other times, for instance, when you have to decide whose side to take in a case of disagreement, it can involve giving more credibility to your friend's point of view, than to that of a non-friend, even though you don't possess good reasons to doubt that the non-friend is telling the truth.

Now, when we start to talk about regulating our judgements and beliefs for reasons of friendship, another, more philosophical, worry arises. Is epistemic partiality just a mild bias—a matter of viewing their character through rose-tinted spectacles—or does it lead you to form and

hold beliefs that are downright irrational—such as if your judgements about them were completely unsupported by your evidence? Before we get ahead of ourselves, though, we have to ask whether popular wisdom is correct when it comes to the benefit of the doubt in friendship. In other words, is it really the case that you should ever give your friends the benefit of the doubt in any of the ways described above? That concern will become the central question of this thesis.

The project

Should we treat our friends with epistemic partiality? That is the central research question of this thesis. This question will be answered affirmatively in the second chapter, where I'll argue that it is indeed the case that you should treat your friends with epistemic partiality, albeit within principled limits. In particular, I'll argue that we have *pro tanto* moral reasons to treat our friends with epistemic partiality, such that it's often the rational thing to do, all things considered. Before we can properly get to grips with the primary research question, though, a preliminary question must be answered. First, what does it mean to treat someone with epistemic partiality? That is, what kinds of behaviours and thought-processes does the practice of treating someone with epistemic partiality involve? Then, we can consider and, indeed, respond to, some epistemological and ethical objections that arise.

Why epistemic partiality in friendship?

There are three points of particular philosophical interest that are developed in this study of the epistemology and ethics of epistemic partiality in friendship.

First, this thesis presents an important new finding about the nature and demands of close friendships and, plausibly, other close, personal relationships. In particular, this thesis attempts to provide some philosophical justification of the admittedly controversial claim that we should treat our friends with epistemic partiality. Building on the Aristotelian insight that friendship is a caring relationship, and in the light of recent work in social epistemology about harmful beliefs, this thesis argues that, if someone is one of our friends, we should treat them with epistemic partiality, at least insofar as doing so helps us to protect their well-being. What's more, once we understand how the fact that friendship is a caring relationship entails that we should treat our friends with epistemic partiality, we can more easily identify other situations in which we ought to treat others with epistemic partiality.

These findings relate to a broader debate in social epistemology about beliefs that harm. This debate took a great step forward in 2018 with the publication of a special edition of *Philosophical Topics* entitled ‘When Beliefs Wrong’, edited by Rima Basu. This volume comprises several papers which argue that beliefs can be harmful, including contributions from Nomy Arpaly and Anna Brinkerhoff (2018), Basu (2018), Simon Keller (2018), Sarah Paul and Jennifer Morton (2018), and Mark Schroeder (2018). The project developed in this thesis shows that we can fruitfully integrate the findings of this new area of research into research on topics that form the core of the discipline.

Second, this thesis presents a new way of understanding the epistemological character of epistemic partiality. While this thesis focusses on the case of friendship, the characterisation of epistemic partiality presented herein is of more general philosophical interest. It’s important to understand the capacities we have for shaping our judgements and beliefs, for instance, by gathering evidence in biased ways or by withholding judgement at opportune moments. Moreover, this thesis also offers (though stops short of endorsing) a principled way to understand the epistemic (ir)rationality of the beliefs that we form as a result of epistemic partiality. In particular, attention is drawn to the relevance of our epistemic goals to questions about when and to what extent we might be morally or prudentially required to act in epistemically irrational ways. This is valuable not only for elucidating the phenomenon of epistemic partiality in friendship, but also presents a possible way forward in our thinking about other occasions of epistemic partiality and belief-regulation more generally.

Third, this thesis helps uncover a tension between moral, pragmatic, and epistemic norms. The case of epistemic partiality in friendship gives us reason to believe that there is indeed an open conflict between pragmatic and epistemic norms. The upshot of this is that sometimes it’s not prudentially rational for you to be epistemically responsible, and other times, it’s prudentially rational for you to be epistemically irresponsible. This fits into a broader debate within epistemology about the nature of epistemic justification and its relation to pragmatic justification. Stroud (2006) first proposes that the phenomenon of epistemic partiality in friendship could be evidence that evidentialism isn’t a tenable theory of epistemic justification. The spirit of Stroud’s suggestion can be found in the defences of several, radical views about the nature of epistemic justification and evidence. Jason Kawall (2013) and Schroeder (2018a) suggest that the phenomenon of epistemic partiality in friendship, and in particular, the fact that it seems rationally permissible in some cases but not in others, corresponding with the stakes, is evidence that pragmatic encroachment-type views of epistemic rationality are correct. Katherine Hawley (2014), similarly, suggests that the phenomenon of epistemic partiality in friendship

could be evidence that permissivism about evidence is true. I argue that we should be reluctant to embrace such radical responses to phenomena such as the phenomenon of epistemic partiality in friendship.

Part One: We should treat our friends with limited epistemic partiality

In the first section of this thesis, I present a theory of epistemic partiality in friendship. This section has two main objectives, each of which is addressed in one of the first two chapters. The first chapter focusses on identifying the phenomenon of epistemic partiality ‘in the wild,’ so to speak, by considering a variety of literary and philosophical sources. Two distinctive practices are identified. The second and perhaps most important chapter of this thesis argues that we should treat our friends with epistemic partiality of the type described in the preceding chapter.

Chapter One: Epistemic partiality in friendship

In this chapter, I introduce the phenomenon of epistemic partiality in friendship by considering a number of examples. The purpose of this is to bring into the light some of the characteristic behaviours by which we manage the judgements and beliefs we form about our friends. I’ll then build an account of these epistemic behaviours to model how the biases with which we seem to treat our friends constitute a distinctly epistemic form of epistemic partiality.

When we consider a friend’s shortcomings of character and conduct, it’s common enough to talk about giving them the benefit of the doubt. After all, so the thought goes, we owe it to our friends to cut them some slack before we form very critical judgements about them. It’s better to be safe than sorry when it comes to judging whether to believe that a friend has disgraced themselves, and the real danger, at least in some cases, is the danger of harming them by thinking poorly of them. It’s not just when we consider a friend’s possible shortcomings that we treat them with this kind of bias. It’s not uncommon to be extra attentive to your friend’s virtues. Indeed, it seems as if we expect our friends to take a more favourable view of our character and conduct than a non-friend might. The point is that there’s some intuitive appeal to the idea that we treat our friends differently from the way we treat non-friends, in the way that we manage our beliefs about them.

I start by considering a case in which it seems someone fails to treat a loved one with the (epistemic) partiality they’re due, given their relationship. This is the case of Othello. I argue that one of the ways in which Othello appears to wrong Desdemona is in his failure to give her the

benefit of the doubt. Rather than withhold judgement or seek to establish her innocence, Othello is swayed by the evidence presented to him by Iago. Consideration of this case also motivates the research question answered in Chapter Two, that is, why, if at all, should we treat friends with epistemic partiality?

Then, reflecting on Carol Reed's film *The Third Man* (1949), and considering a variety of examples drawn from Keller (2004, 2018), Stroud (2006), and Hazlett (2017), two practices of epistemic partiality are identified. First, it seems we treat our friends with epistemic partiality in the way that we gather evidence about them. Second, it seems we treat our friends with epistemic partiality in the way that we process the evidence we possess.

Finally, I consider some of the limits of epistemic partiality. In particular, I focus on cases in which it's not clear that someone would or should treat a friend with any kind of epistemic partiality. These are discussed in more detail in later chapters of the thesis.

Chapter Two: Why we should treat our friends with limited epistemic partiality

At this stage, we've identified the key behaviours of epistemic partiality in friendship and we've seen that that these behaviours involve the violation of some plausible epistemic norms of impartiality. But are we actually required to treat our friends with epistemic partiality, or is it just a bad habit that doesn't contribute to our relationships? In this chapter, I argue that we should sometimes treat our friends with epistemic partiality.

The argument I propose is called 'the argument from care'. Friendship, according to just about any account of friendship on offer, is a caring relationship. If someone's your friend, then you care about them. The nature of care is no doubt a complicated matter, involving doxastic, affective, and volitional states, including desire. For the purposes of this argument, though, I focus on desire. On this view, if you care for someone, then you desire their well-being and, what's more, you desire their well-being for its own sake. What has this got to do with epistemic partiality?

Well, it's plausible that your beliefs and judgements about someone can harm them, but it's not clear exactly how. Rather than defend a particular theory of harmful beliefs, I defend a broad account of the phenomenon. First, following Keller's lead (2018), I explain how, on any mainstream understanding of well-being, it's plausible that beliefs can harm. Then, I explain three hypotheses about how beliefs can harm, and I apply them to the case of friendship. First, I argue that beliefs which express undue criticism can harm, even if that criticism is epistemically justified from the point of view of the believer. Marušić and White (2018) defend an

epistemology of harmful beliefs, drawing on Strawson (1962), which explains that your beliefs can harm another person when they fail to manifest an appropriate amount of goodwill towards that person. After reflecting on these hypotheses and how they relate to the cases described in Chapter One, I conclude that there are good reasons to think that there are *at least* two different ways that beliefs can harm our friends. Then I consider an account of harmful belief which is due to Basu (2019). Basu argues that beliefs can harm, in particular, when they fail to accommodate people's deeply held desires about how others relate to them, in particular, with relation to their practical identity involved. Each of these hypotheses can explain how our beliefs and judgements might harm our friends. Furthermore, our friends are especially vulnerable to such harms because of the expectations they have of us, as their friends.

Given that (i) we desire our friends' well-being, and (ii) our beliefs and judgements about them can be inimical to their well-being, it's reasonable to think that we have some *pro tanto* moral reasons to treat our friends with epistemic partiality, or so I argue in the final stage of the argument from care. That's because treating your friend with epistemic partiality is the best way to protect their well-being from the harm that your beliefs about them can inflict. Considering a variety of counter-examples, that is, cases when it seems we shouldn't treat our friends with epistemic partiality, I finally settle on the conclusion that we should treat our friends with limited epistemic partiality. I'll also discuss cases of conflicting loyalties. What should you do when one friend, A, tells you about the gross wrongdoing of another friend, B? I'll attempt to provide a principled explanation of when you should treat B with epistemic partiality.

Part Two: Responses to some objections

Having presented a theory of epistemic partiality in the first part of this thesis, the second part comprises responses to a number of responses to the claim that we should treat our friends with limited epistemic partiality. I'll briefly describe the spirit of these objections, before summarising each chapter in turn.

Chapter Three responds to an epistemological objection. According to the epistemological objection, we shouldn't treat our friends with epistemic partiality because it would be epistemically irrational to do so, and these epistemic reasons override any moral reasons we might have for treating our friends with such partiality. Chapter Four responds to a moral objection to the claim that we should treat friends with any amount of epistemic partiality. If we treat our friends with any amount of epistemic partiality, we risk insulting non-friends. To make matters worse, given some plausible assumptions about how we come to make friends,

treating them with any amount of epistemic partiality is likely to perpetuate epistemic injustice. But these moral consequences are intolerable, and therefore, it's not the case you should treat your friends with any amount of epistemic partiality. Finally, Chapter Five responds to an objection from within the ethics of friendship. It's constitutive of friendship that what beliefs you hold about your friends are based in the facts about them, and that they're not, for example, held for reasons of friendship. But the claim that we should treat our friend with any amount of epistemic partiality implies that we should hold beliefs about our friends which are not based in the facts but rather are held for reasons of friendship. Therefore, so the objection goes, it's not the case that we should treat our friends with any amount of epistemic injustice.

Chapter Three: An epistemological objection to treating friends with epistemic partiality

Now that we've got a working model of the behaviours that constitute epistemic partiality in friendship, a question naturally arises regarding whether, and to what extent, those behaviours are epistemically irrational. In part, this chapter can be seen as a response to some objections levelled by Arpaly and Brinkerhoff (2018).

Is epistemic partiality in friendship epistemically irrational? First, I explain why supporters of the epistemological objection might think that it is. Given that epistemic partiality involves, on my view, partiality at the level of evidence gathering and evidence processing, I provide arguments in favour of thinking that epistemic partiality is epistemically irrational at each of those levels. So why would the objector think that epistemic partiality in evidence gathering is epistemically irrational? The thought is that to the extent that what it's epistemically rational for us to believe about someone is determined by our evidence, our personal relationship to them makes no difference to how you should gather evidence about them. Though this principle appears to be strictly moral in character, I explain how it may follow from our epistemic goals, namely, believing the truth and avoiding error (cf. William James 1887/1992, also Alston 1985, Bonjour 1985).

Second, I explain why the objector might think that epistemic partiality in evidence processing is epistemically irrational. The thought is that, to the extent that what it's epistemically rational for us to believe about our friends is determined by our epistemic standards, our personal relationships make no difference to which epistemic standards are the correct standards for us to apply. In defence of this principle, I argue that the facts about a relationship make no

difference to how much weight it's epistemically rational to give particular pieces of evidence in our reasoning.

In the remainder of the chapter, I respond to the epistemological objection. First, I explain why there may be good reasons for thinking that epistemic partiality in friendship isn't epistemically irrational after all. Regarding epistemic partiality at the level of evidence gathering, I consider the arguments of Jason Kawall (2013), Sanford Goldberg (2018), who argue that the processes of epistemic partiality are not in fact reasonable acts of caution in response to the risk of error. Second, I consider two responses to the claim that epistemic partiality in evidence processing is epistemically irrational, from Katherine Hawley (2014) and Schroeder (2018). Hawley argues that if permissivism is true, then it's not epistemically irrational to apply different standards when forming beliefs and judgements about our friends. Schroeder argues that many cases of epistemic partiality in evidence processing are best understood through the lens of pragmatic intellectualism. According to Schroeder's response, it's epistemically rational to withhold judgment when the risks of error are high. Friendship is one such case.

For the sake of argument, I remain agnostic about whether epistemic partiality is epistemically irrational. Instead, I continue to argue that even if epistemic partiality is epistemically irrational, there are good reasons to reject the claim that this implies that we shouldn't treat our friends with even limited epistemic partiality. Developing Keller's (2004) insight, I argue that the epistemological assumption assumes that epistemic and moral reasons are commensurable. This assumption is dubious, or so I argue. Finally, I argue that even if the dubious assumption is correct, then it still doesn't follow that our *pro tanto* moral reasons to treat our friends with epistemic partiality are overridden by our epistemic reasons not to.

Chapter Four: A moral objection to treating friends with epistemic partiality

We should treat our friends with limited epistemic partiality, or so I argue in this thesis. But when we commit to treating our friends with epistemic partiality, we risk insulting non-friends, particularly, when we refuse their testimony because of the harm it would inflict on our friends if we accepted it. It's bad to insult people, but arguably, epistemic partiality in friendship is more problematic yet. You might worry that by treating friends' testimony with epistemic partiality we risk committing and perpetuating epistemic injustice.

In this chapter, I consider the argument, first proposed by Sheila Lintott (2015), that epistemic partiality in friendship is morally problematic. In its strongest form, this argument can be considered an objection to the thesis that friendship requires epistemic partiality at all. To

begin, I explain the nature of the special insult of refusing someone's testimony and how we risk paying such an insult to non-friends when we refuse their testimony as a result of our epistemically partial treatment of our friends. In this part of the discussion I reflect on recent work on testimonial insult by Allen Hazlett (2017) and Finlay Malcolm (2018a, 2018b). Then, building on a theory of testimonial silencing defended by Kristie Dotson (2011), I explain why you might think that epistemic partiality in friendship might not only cause insult, but also, epistemic injustice. Drawing again on Lintott (2015), I explain why this might be taken to present a serious challenge to the tenability of the thesis that friendship requires epistemic partiality. Particularly, I consider how a kind of homophily, that is, the tendency to befriend people with similar characteristics as oneself, coupled with epistemic partiality in friendship, might lead to epistemic injustice. To borrow Lintott's example, if we're more likely to befriend people of the same ethnicity, and treat those people with a positive epistemic bias, we're therefore likely to be treat non-friends who aren't of the same ethnicity with a negative epistemic bias. This, of course, would be a morally abhorrent implication of the view that I defend in this thesis.

Ultimately, though, I argue that while friendship requires limited epistemic partiality, it doesn't require the kind of epistemic partiality that would constitute or perpetuate epistemic injustice. Developing the work of Jessica Isserow (2018), I explain how our moral priorities interact with the duties we have towards our friends. The importance of our friends' well-being occupies a place in our moral priorities somewhere between the importance of not insulting strangers and not committing or perpetuating epistemic, in this case testimonial, injustice.

Chapter Five: An objection to treating friends with epistemic partiality from the nature of friendship

In the fifth and final chapter of this thesis I consider one last objection to the theory of epistemic partiality defended in part one. Lindsay Crawford (forthcoming) presents an argument from friendship against epistemic partiality in friendship. She argues that it's constitutive of friendship that the judgements we form about our friends be responsive to the facts about them. If our beliefs about our friends are based on state-given reasons (i.e., the fact that we're friends) rather than object-given reasons (the facts about our friends), then, it seems, we've failed to appreciate our friends for who they really are—we might even worry that we don't know them properly. But believing for state-given reasons, Crawford argues, is exactly what those who defend epistemic partiality in friendship are promoting. Therefore, she concludes, the partialists must be mistaken: friendship doesn't require epistemic partiality.

I don't rebut Crawford's argument right away. Rather, I develop the concerns presented by Crawford through a close consideration of the properties view of love, according to which the correct reasons for loving someone are the properties they possess (Neil Delaney 1996, Keller 2000, Bennet Helm 2010, Sara Protasi 2016). If we accept the properties view of love, I argue, then Crawford's case against epistemic partiality in friendship appears even stronger.

My rebuttal comes in two parts. In the first, I draw attention to the difference between the Crawford's conception of epistemic partiality in friendship and my own. Crawford's presentation of the character of epistemic partiality in friendship is rather stronger than the view I've defended in this thesis. Where she argues that epistemic partiality in friendship requires believing for state-given, pragmatic reasons, I've argued that epistemic partiality may also include behaviours such as: (1) gathering evidence in biased ways; (2) processing your evidence in biased ways; (3) updating your evidence in biased ways. None of these forms of epistemic partiality necessarily involve believing for state-given reasons. State-given, pragmatic reasons appear to play a different role.

In the second, I argue that Crawford's account of friendship is mistaken. It's not the case, I argue, that friendship between A and B excludes A's believing that p , when p is about B, for state-given reasons. First, I argue that sometimes, we deliberately petition our friends to believe or withhold belief for state-given reasons. Second, I argue that we sometimes consider other people to be good friends, because they are inclined to believe things for state-given, rather than object-given, reasons. Finally, reflecting on some remarks by C.S. Lewis (1960) in *The Four Loves*, I explain why I think we shouldn't be alarmed by the thought that friendship involves incomplete or partial knowledge of our friends' characters.

Conclusion

To conclude, I'll propose a further line of research. I've already argued that in situations where our beliefs harm, we must consider whether we should treat the victim of that harm with epistemic partiality. If we fail to accept the reports of victims of sexual or racial discrimination or violence, we risk causing them serious harm—compounding the damage that has already been done. Should we treat the self-reporting victims of sexual or racial discrimination or violence with epistemic partiality when they report the crimes that have been committed against them? I'll discuss this at greater length in the conclusion of this thesis.

In this thesis, I argue that we should treat our friends with limited epistemic partiality. Friendship is a caring relationship, and as such, when our friends' well-being is at stake, we

should to take appropriate measures to protect it. When our forming a given belief or judgement about a friend will harm them, we should treat them with epistemic partiality. But just what is it, to treat someone with epistemic partiality? That's the question that I'll answer in Chapter One.

Chapter One

Epistemic partiality in friendship

§1 Introduction

In this chapter, I'll describe how we show friends epistemic partiality. In Section 2, I'll present a range of cases, largely drawn from literary sources. In the following two sections, I'll characterize just what showing a friend epistemic partiality is by dividing the phenomenon into two parts: one which concerns evidence gathering, and another which concerns evidence processing.

In Section 3, I'll argue that the fact someone is a friend sometimes guides the way we gather evidence about them. Where our friends are concerned, we look for evidence of the goodness of their character and conduct, and we don't seek out evidence of their failings: sometimes we even avoid it. We don't normally root around for proof of our friends' moral failings or social shortcomings, and it's not obvious that we're even-handed when we do look for evidence. Rather, we give them the benefit of the doubt when they face criticism, and we look on the bright side when we consider their conduct and how it reflects their character.

In Section 4, I'll argue that the fact that someone is a friend can also shape the way we process the evidence we possess about them. We might require much more evidence to believe something bad about a friend, and much less evidence to believe something good about them, than we would require to form the same beliefs about a stranger. Being someone's friend also makes a difference to whether we accept their testimony. We're inclined to believe what our friends tell us, perhaps even when a reasonable non-friend wouldn't accept their testimony. Conversely, we're likely to ignore what people tell us about our friends, if accepting what they tell us would make us see our friends in a less favourable light. Thus it's not only in how we go about gathering evidence that we show bias towards our friends, but also in what we do with the evidence we've got.

We can think of this bias as a kind of epistemic partiality because, while we treat our friends in this way, we don't extend the same favour to non-friends. It's important to note how different this epistemic behaviour is from the behaviour we'd expect from a disinterested third party. Disinterested third parties don't give each other the benefit of the doubt when they judge each other's conduct or character. Plausibly, if someone is judging the performance of a colleague with whom they share no close bond, then they're unlikely to expend extra effort looking for the goodness in their colleague's work and avoiding evidence of their failings. They

won't set a low threshold for evidence for belief in their colleague's goodness. In summary, as a disinterested third party, they're less likely to reach a biased judgement about their colleague than, say, one of their colleague's friends.

But before we discuss the epistemology of epistemic partiality in friendship, it'll be helpful to consider how our beliefs can harm and benefit other people. In Section 2, I'll start by presenting a case from which epistemic partiality is lacking, and where its absence has deeply regrettable consequences. This is the case of Othello and Desdemona.

§2 Beliefs and relationships

In Shakespeare's eponymous tragedy, Othello is so gravely misled by rumours of the infidelity of his wife, Desdemona, that he murders her in a fit of jealous rage. Later, upon realising his mistake, he commits suicide. Othello and Desdemona are the victims of a conspiracy contrived by Iago and Rodrigo. Iago is one of Othello's trusted advisors, but he hates Othello for promoting another man, Cassio, instead of him. Rodrigo hates Othello because he, Rodrigo, asked Desdemona's father for her hand in marriage, only to find out later that she has eloped with Othello. The turning point in the play takes place in Act III, Scene iii, when Iago convinces Othello that Desdemona has been sleeping with Cassio. Initially Othello rejects the very suggestion that his wife would be unfaithful: 'Nor from mine own weak merits will I draw The smallest fear or doubt of her revolt; For she had eyes, and chose me. No, Iago, I'll see before I doubt, prove' (Act III, Scene iii, 201-203). Othello is, at this stage, resolute in his disbelief of the rumours around Desdemona's conduct, and steadfast in his confidence in her faithfulness.

This does not last. After being strung along by Iago's insinuations, Othello starts to have doubts about Desdemona. Now worried, he says to Iago:

I think my wife be honest and think she is not;
I think that thou art just and think thou art not.
I'll have some proof. Her name, that was as fresh
As Dian's visage, is now begrimed and black
As mine own face. If there be cords, or knives,
Poison, or fire, or suffocating streams,
I'll not endure it. Would I were satisfied! (Act III, scene iii, 385-90)

Once his suspicions of Desdemona's infidelity are aroused, Othello demands to know more. He says to Iago (Act III, Scene iii): 'Make me to see't: or, at the least, so prove it, That the probation

bear no hinge nor loop 'To hang a doubt on; or woe upon thy life!' In asking Iago for proof of Desdemona's infidelity, Othello shows that he no longer believes that Desdemona was faithful to him, at least, not with any great confidence. Indeed, he seems to have suspended his judgement on the matter. Othello and Desdemona's problems have only just started, though. Once Othello's eyes are opened to the possibility that Desdemona is unfaithful to him, his suspicions start to grow and grow. This stark change in his judgement about her character and conduct is reflected in his changing behaviour, particularly in his decision to send Iago to gather evidence about Desdemona's conduct. To put an epistemological spin on this, we might say that Othello seeks more evidence when he suspects that Desdemona is unfaithful, but he lacks sufficient evidence to believe it with any confidence.

What should we make of this? You might think that this is what epistemic propriety demands of us. Indeed, that's the correct way to think about this case, at least according to the epistemological objection I'll defend in Chapter Three. According to the objection, this is just a matter of common sense: If you're going to make a judgement about whether p , you should proportion your belief to your evidence. If you don't have much evidence and you don't want to suspend your judgement, then you should look for more evidence.¹ We might think, though, that Othello owes Desdemona even more than this. It's not quite right that he ought, if not coolly, at least impartially, to gather and process his evidence. As Stroud writes, it seems that we 'owe [our] friends something other than an impartial and disinterested review of the evidence where they are concerned' (2006, 504). In any case, there isn't enough evidence to justify Othello's belief that Desdemona is unfaithful. Moreover, it seems wrong that he's willing to look for more evidence to settle the matter. But this is exactly what Othello does, and with tragic consequences.

It's hard to doubt that Othello and Desdemona are both victims, albeit of different sorts, in this tragedy.² And for all this, it seems that Othello has failed in his duties as Desdemona's

¹ Robert Hall and Charles Johnson (1998) defend this principle of epistemic rationality as the epistemic duty to gather more evidence. According to Hall and Johnson, 'for any proposition that is less than certain on one's present evidence, one has an epistemic duty to seek more evidence about that proposition' (1998, 133).

² There are many ways to read Shakespeare's *Othello*. Some read it as a play about pathological jealousy, others as concerning the consequences of internalised racism. The reading I present here is not intended to compete with or cast judgement on these readings. My only aim in selecting the tragedy of Othello is to illuminate the phenomenon of epistemic partiality in close, personal relationships by presenting a notable example of its absence.

husband. He comes to believe that she's unfaithful based on rumour and insinuation (and, in the end, on the reported discovery of her handkerchief in Cassio's lodgings.) In Act III, we see how Iago manipulates Othello's jealousy and eventually sets him down the path which ends with him convinced of Desdemona's infidelity.

How does all this affect Desdemona? Even before Othello murders her in Act V, scene ii, she suffers immensely as a result of his jealousy. Desdemona suffers physical violence when Othello strikes her (Act IV, scene i, 243) and verbal violence when he taunts her, saying 'I cry you mercy, then. I took you for that cunning whore of Venice that married with Othello' (Act IV, scene ii, 92). Beyond the physical and verbal violence she suffers, she's also deeply insulted, I suggest, by Othello's beliefs. Desdemona says to Emilia, her attendant:

Unkindness may do much;
And his unkindness may defeat my life,
But never taint my love. I cannot say 'whore.'
It does abhor me now I speak the word;
To do the act that might the addition earn
Not the world's mass of vanity could make me. (Act IV, scene ii, 166-171)

That Othello's beliefs impugn her good character causes Desdemona a great deal of distress. His beliefs seem to be a distinctive source of harm for her, in addition to the verbal and physical violence that he inflicts on her.

We might be able to draw a more general lesson from this example about the relationship between beliefs, judgements, and harm.³ Perhaps what you believe about someone can harm them, especially when your beliefs impugn their moral character or their conduct. That said, the case under discussion is extreme. It's extreme because what Othello comes to believe about Desdemona is that she's betrayed him in such a way that explicitly warrants the termination of their relationship. To sleep with someone other than your spouse, we can suppose, is to cross a line that cannot ordinarily be uncrossed. This could be cashed out as a matter of betrayed trust. The supposed betrayal is so extreme that Othello can no longer trust Desdemona.

³ The relationship between beliefs and harm will be discussed in greater length in the following chapter, with particular reference made to the work of Basu (2019), Basu and Mark Schroeder (2019), Simon Keller (2018), Berislav Marušić & Stephen White (2018), and Mark Schroeder (2018a).

Let's dwell for a moment on the subject of trust. We might ask whether there is (or should be) epistemic partiality in trusting our friends? This is a tricky point that depends largely on what account of trust you accept. If trust is doxastic, that is, if trusting someone to F involves believing that they will F, then it's at least possible that trust involves epistemic partiality. But, as Katherine Hawley (2014, 2034) argues, if trust is non-doxastic, that is, if trusting someone to F doesn't necessarily involve believing that they will F, then it's likely not the case that trusting friends involves epistemic partiality. In fact, invoking Richard Holton's (1994) non-doxastic account of trust, Hawley claims that it is indeed possible to trust someone without believing that they're trustworthy or that your trust in them won't be frustrated. That's because reliance, not belief, is central to the attitude of trusting, or so the thought goes.

It may well be the case that trust is a non-doxastic attitude and that there is no epistemic partiality in trusting. Plausibly, though, this wouldn't diminish the importance of *believing* that a friend is trustworthy very much. Even if your trusting someone doesn't require you to believe that they'll do what they promised you, it seems plausible that your trusting them is incompatible with your being doubtful or even disbelieving that they will do what they promised you. That explains what sounds unnatural about the sentence 'I trust Adam will keep his promise, but I don't believe that he'll keep his promise.'⁴ I'll explore this in greater detail in the following chapters.

In the case of Othello and Desdemona, it seems as if some beliefs and judgements can harm our friends and close, personal relations. Plausibly, our beliefs and judgements can benefit our friends too.⁵ A good friend can be like a coach, in that they seek to help their friends improve themselves and achieve their goals, according to Keller (2004, 339-340). Being a coach has an epistemic component which can sometimes require a kind of reassuring bias.⁶ Consider this passage from Keller:

Competitive running is not centrally concerned with seeking the truth. Among the things I want from my running friends is feedback that will keep me motivated. Sometimes that purpose is served by a completely accurate assessment of my performance and my prospects. But when I

⁴ This seems like an example of Moore's paradox. Whether it's actually paradoxical depends on what account of trust you accept, among other things.

⁵ This applies to the other doxastic attitudes, disbelief and withholding of judgement too.

⁶ Contra objections from Jason Kawall (2013, 355), Keller (2018, 28) argues that there's no reason to think that your showing a positive bias in favour of your friends is inimical to your role in their self-improvement.

am fighting against age and injury and poor form, sometimes I just need to hear something positive, and it is valuable to have friends who see the best in me and believe the best about my potential to improve. I know that as my friends, their opinions are biased and I should interpret their feedback in that context. And I know that if I were a stronger, less needy person, then I would have no interest in being encouraged and supported; I could take care of that myself. But still, being the person I am, those are some of the good things I get from having friends. (Keller 2018, 27)

The thought is that, sometimes, as a good friend to someone, you'll be quick to notice their strengths and you'll be resistant to believing that they have any grave flaws. You'll do this for the sake of promoting their self-esteem or some other aspect of their well-being. Thus it's not uncommon to believe 'for someone else's sake', to borrow Keller's phrase (2018). That seems like a central element of epistemic partiality. One aspect of what Keller is suggesting is that we show a kind of leniency to our friends, in that we demand less evidence to believe something positive about them than we would of someone with whom we didn't share a close, personal relationship. When looking on the bright side, so to speak, epistemic partiality may involve overweighing evidence which supports a beneficial judgement about our friends. When giving the benefit of the doubt, epistemic partiality can involve ignoring or underweighing evidence which would cast doubt on our positive judgements about a friend.

Our beliefs and judgements, so it seems, can harm our friends. But if we could control our beliefs, we would be able to confer positive and negative benefits upon our friends: They confer negative benefit when our beliefs prevent a harm that they would otherwise suffer. They confer positive benefit when our beliefs confer some good that wouldn't otherwise be conferred. In fact, this seems like something we actually do, namely, when we treat our friends with epistemic partiality. In the remainder of this chapter, then, I'll consider what it means to treat our friends with epistemic partiality, first, in the way that we gather evidence about them, second, in the way we process that evidence, and third, in the way that we update our belief based on our processed evidence.

§3 Evidence gathering

The Third Man, Carol Reed's 1949 film, relates the thrilling story of Holly Martins' efforts to track down his old friend, Harry Lime, who has disappeared under mysterious circumstances in the

ruined city of Vienna.⁷ Martins, a penniless writer of pulp fiction, arrives in Vienna, where he plans to meet Lime. They were close friends before the war, but they have had little contact since fighting began in Europe. A few years before the war, Lime offered Martins a job in Vienna, and at long last Martins has come to accept his offer. Unfortunately, Martins' trip to Vienna gets off to a bad start. No sooner has he arrived than he finds himself at a funeral—Lime's funeral. It's not long before Martins becomes suspicious about the circumstances around Lime's last days and death. He is troubled by the inconsistency among the various accounts of Lime's death that he receives from the people who knew him. He resolves to find out the truth about his friend's disappearance.

Martins soon encounters Major Calloway, a member of the British military police stationed in Vienna. (At the time, just after the end of the Second World War, the administration and policing of Vienna was shared between British, French, American, and Russian military forces.) Calloway urges Martins to stop his investigations and to leave Vienna immediately. It becomes apparent that Lime is suspected of involvement in some kind of criminality. Given what he knows about Vienna at that time, Martins assumes—perhaps reasonably—that Lime is suspected of trading on the black market. When Calloway urges Martins to stop looking for Lime and to leave Vienna, Martins is astonished. The fact that Calloway wants Martins to leave Vienna only makes Martins believe that there is a conspiracy afoot against Lime. Why else, he thinks, would Calloway try to send me away from Vienna?

This all changes in a dramatic sequence in which Calloway reveals the true, ghastly nature of the allegations against Lime to Martins. It turns out Lime is not suspected of some petty crime like smuggling cigarettes. Rather, Lime is suspected of stealing penicillin, diluting it, and selling it back to hospitals. Many people die as a result of his crime, and those who survive live on in agony. When Martins is given a tour of a children's hospital where many of Lime's victims are apparently bound to remain until their deaths, he cannot be mistaken any longer—the veil is lifted. The nature of the allegation against Lime is utterly awful and there's little doubt of his guilt.

This film is interesting from an epistemological perspective because it illustrates how we shape the way we gather evidence in order to protect our friends' well-being. As a viewer we find ourselves in the position of an impartial, third party, while we watch Martins' initial resistance to, and then acceptance of, the evidence against Lime. In a sense, the viewer shares their evidence

⁷ Anthony Cross suggests this as an example of epistemic partiality in his manuscript, "The puzzle of partiality".

with Martins. The viewer is entitled to take it for granted that before the War, Lime was a mischievous rascal but not a murderous criminal mastermind. It becomes clear quite quickly to the audience that Lime has changed. He's no longer the man whom Martins cherished as a friend back in the years before the War. We know this because we witness numerous reports from people who either knew Lime well, or who have no obvious reason to hide the truth.⁸

Martins does what he can to avoid evidence of Lime's wrongdoing. One of the main methods that Martins employs is to exclude certain pieces of evidence. When Major Calloway tells him that Lime is a criminal, Martins doesn't even recognise Calloway's utterances as genuine attempts to tell him anything. Rather, he judges them to be attempts to waste his time and distract him. From Martins' perspective, Calloway's protests are more like bullshit than lies—he doesn't think that Calloway is trying to get him to believe something false, but rather that Calloway is trying to lead him to believe whatever it takes to make him go away, with little concern for the truth of those beliefs.

It's important to emphasise that Martins' friendship with Lime makes a difference to the way he gathers evidence about him. Martins tries to gather evidence which will prove Lime's innocence, and he deliberately refrains from gathering evidence of his guilt. But we should consider how this differs from the behaviour we would expect from a disinterested third party. A disinterested third party, that is, someone who had nothing to gain or lose by judging Lime negatively, is unlikely to feel pressured to treat him with bias by any sense of commitment or loyalty. What's more, for a third party, the stakes are inverted: It's possible that the priorities of the average citizen of Reed's Vienna has an interest in uncovering and convicting the perpetrators of crimes such as those alleged to have been committed by Lime. So perhaps, the ordinary citizen, unlike Martins, would strive to gather representative evidence.

This is a very serious example of how we might gather evidence in a biased manner in order to protect a friendship. Now let's consider a more everyday example from Keller (2004). I'll refer to this as Keller's Poetry Reading example hereafter.

Rebecca is scheduled to give a poetry reading at a café. [...] She lets her good friend Eric know that she'll be giving the reading, and asks whether he'd mind coming along to be the audience. [...] [I] think about Eric's beliefs after Rebecca's performance has concluded. If her poetry really is brilliant, then he'll probably believe that it is. If her poetry is truly awful, then he'll probably

⁸ Granted, this is a thriller and there's some suggestion of a conspiracy against Lime. But mainly, I think, this suggestion comes from Martins' own rationalisation. It doesn't seem to me that viewer is ever given a strong reason to think that Lime's the victim.

notice. If Eric behaves as a good friend would, however, then there will be possible situations under which Eric will believe that Rebecca's poetry was pretty good, and that there's a decent chance that the literary agent will show some interest in publishing it, even though he would not have those beliefs about the work of a stranger who read exactly the same poem. (Keller 2004, 331-333)

We can imagine that Eric might focus his efforts on looking for signs of talent in Rebecca's poetry. He'll not look out for weaknesses and he'll purposefully avoid paying too much attention to any shortcomings that he notices. That's not all Eric could do. Imagine that Derek, a friend of both Eric and Rebecca, learns that Rebecca gave a poetry reading and that Eric attended. Derek has heard that the venue has a reputation for somewhat substandard poetry and this strikes him as odd, because he always assumed that Rebecca was a good poet. So, to protect his positive appraisal of Rebecca's poetry, he asks Eric about it next time he sees him. 'I bet Rebecca's performance was great, wasn't it? Isn't she getting better every time?' Of course, by asking a leading question like this Derek implicates that he doesn't want to hear negative criticism of Rebecca's reading. Rather, Derek invites Eric to share praise about Rebecca's reading.

Thus 'the fact that a person is someone's friend can explain why that person is inclined to believe certain sorts of falsehoods', or so writes Simon Keller (2004, 330). Well, perhaps Eric is at risk of believing certain sorts of falsehoods about Rebecca's poetry reading. In particular, he might come to hold a more favourable judgement of her ability than he would if they were not friends.

How, then, does gathering evidence in epistemically partial ways allow us to form more favourable judgements about our friends? Well, we can identify two distinctive behaviours in this variety of epistemic partiality. In some cases of epistemic partiality at the level of evidence gathering, we seek dis-confirmatory evidence towards certain hypotheses about our friends, in particular, if our believing those hypotheses would harm them. This is what we see in the *The Third Man*, at least before Martins learns the true, moral gravity of the accusations against Lime. He seeks to discredit what evidence he possesses of Lime's guilt and also to find proof of his innocence.

In other cases of epistemic partiality in friendship, we search for evidence that would confirm or support our holding beliefs that benefit our friends. In Keller's Poetry Reading example, Eric searches for evidence which confirms his view that Rebecca is a talented poet. He does this, I suggested, by exclusively looking for signs of talent. He also avoids looking for flaws in her performance. This enables him to construct a case for coming to a favourable judgement

about her poetry, thereby enabling him to confer whatever positive benefit his good opinion has for her well-being.

Epistemic partiality in friendship involves a distinctive form of bias at the level of evidence gathering. This bias helps us to form doxastic attitudes which both positively and negatively benefit our friends. Whether this is epistemically rational will be discussed in Chapter Three, where I'll argue that whether or not epistemic partiality is epistemically irrational, it's nevertheless morally required in our friendships.

§3 Evidence processing

Writing in *The Guardian*, Ijeoma Oluo (2015) explains how she empathises with those who find themselves unable to process the possibility that a close relation could be guilty of treating another with terrible cruelty. Oluo's close relation, whom she refers to as Steve, was twice convicted of sexual assault. It took Oluo and some of her family members twelve years to accept the truth about Steve. She explains how they were able to maintain their belief in his innocence for so long:

When the victim showed up to court covered in cuts and bruises and her neck in a brace, Steve's parents argued to us that it was a ruse: the bruises were makeup; the neck brace was for show. My mom, who hadn't even been particularly close to Steve, believed all of their stories and justifications. She never said anything to vilify the victim, but she would say that something about it just "didn't seem right." (Ijeoma Oluo, 2015)

It seems that by explaining-away evidence through a kind of *post-hoc* rationalisation, Steve's close relations were able to cultivate the belief that he was innocent despite the public availability of overwhelming evidence. They sought to protect his interests and his image in their eyes by biasing their epistemic practices so that they ended up with these beliefs. When we're presented with accusations about the wrongdoing of our friends and loved ones, it seems we demand a great deal of evidence before we believe the accusations against them. While some aspects of the case against a friend may be harder to deny, we can feasibly suspend judgement about certain, key aspects of the case against them. For instance, Steve's close relations suspend judgement about the credibility of the victim's testimony and deny the veracity of her injuries.⁹

⁹ Paul Faulkner (2018) considers a case in which your friend is accused of murder. The question is, should you accept their testimony when they protest their innocence? Faulkner argues that we

It's plausible that they would not have treated Steve in this way, were he not a close, personal relation. Whilst they might not have realised it, they were treating him with epistemic partiality. Of course, in the light of the #MeToo movement and the debacle around the appointment of Bret Kavanaugh to the US Supreme Court, this provokes serious questions about the moral permissibility of epistemic partiality. Other cases of morally-objectionable epistemic partiality in friendship are conceivable. I'll attempt to negotiate this problem in Chapter Four.

The way we are lenient with our friends can also make a difference to how we evaluate the testimony of a friend. To understand this better, I'll explain why the notion of the special insult of refusing testimony. In Chapter Two I'll argue that sometimes we should accept the testimony of our friends because not accepting it risks insulting them. Since insulting our friends harms them, so long as we have a moral reason to protect and promote their well-being from harm, we have a moral reason to accept their testimony.

To illustrate this point, we can return to another fictional example. Allen Hazlett (2017) relates some of the events of David Foster Wallace's 'Oblivion' (2004). In 'Oblivion', Randall Napier reports 'the strange and absurdly frustrating marital conflict between Hope and myself over the issue of my so-called "snoring"' (Wallace 2004, 122). Hope tells Randall that he snores; Randall refuses to accept this and instead insists that she must be dreaming it. His refusal to accept her testimony insults her, and she's insulted, because (from her perspective) he's unreasonably doubtful of her credibility—in particular, that she's incompetent when it comes to distinguishing her dreams from reality. This drives Randall Napier into a kind of sceptical fury: It's not just about snoring for him. It's about his ability to perceive the external world. He reasons thus: He can only be snoring if he's asleep. If he's asleep, then he's unconscious. And if he's unconscious, he doesn't hear Hope's protestations. But he *did* hear Hope's complaints, so he must have been conscious. And since he was conscious, he must have been awake, too! Therefore, it cannot be the case that he was snoring, despite what Hope says. It's unbearable for Randall that Hope casts doubt on his credibility with respect to his own consciousness. Randall understands that Hope is trying to tell him something, but he rejects it because, by his lights, he has stronger evidence against her testimony than she has in support of it.

It damages Hope and Randall's marriage when he refuses her testimony. An irony of this case is that, by the end, Hope and Randall are in very similar situations. They have both told

should, and that it can be epistemically reasonable to do so. I dispute the second part of his position in Chapter Three.

each other what they believe about the ‘snoring situation’, and they both seem to be sincere. And they reject each other’s testimony. This example is amusing but it points to something true: It damages our relationships with people when we do not accept their testimony. Now, of course there must be circumstances in which it’s epistemically impermissible to accept someone’s testimony. In the following chapter I’ll consider why refusing someone’s testimony can constitute a special insult to them.¹⁰

Epistemic partiality at the level of evidence processing isn’t always as simple a matter as refusing or accepting individual, key pieces of evidence as it suits your interests. It’s sometimes a more complicated affair, involving the interpretation of a friend’s conduct. Stroud writes that ‘we will go to greater lengths in the case of a friend to construct and to entertain alternative and less damning interpretations of the reported conduct than we would for a nonfriend’ (2006, 506). She instructs the reader thus:

Suppose, for instance, that a third party reports that your friend Sam recently slept with someone and then cruelly never returned any of that person’s calls, knowingly breaking that person’s heart. (Stroud 2006, 504)

In her discussion of this case, which I’ll return to throughout this thesis, Stroud suggests that we show partiality in the way that we process the evidence we gain from such reports. Sam’s friends might require more evidence before they accept the allegation against him than they would otherwise, for instance, if Sam were not their friend. They might even withhold judgement even in spite of the evidence they possess.

Epistemic partiality in friendship is not always as simple as denying the brute facts about a friend’s behaviour. Sometimes our evidence about the raw facts about a friend’s actions is incontrovertible. But ‘beyond the incontrovertible,’ Stroud writes:

many interpretive avenues open up. And “the proper bias of a close friend” finds a fertile field of operations in this rather murky and uncharted terrain. What is characteristic of the good friend

¹⁰ Admittedly, sometimes it’s possible to ignore the testimony of your close, personal relations, for instance by pretending you didn’t hear them, or more subtly, by pretending to accept their testimony while secretly rejecting it. But this isn’t practicable, at least not all the time. For instance, if someone tells you something directly, several times, and what they tell you demands action that’s only reasonable if you accept what they say, then you cannot pretend that you didn’t hear your friend’s testimony, and you can’t pretend to accept it whilst secretly ignoring it.

seems typically to come in at this level of interpreting the reported actions and placing them in perspective, rather than at the level of denying the base-level facts and events being related by the teller of the story. In general, a good friend is likely to interpret what she hears in a less damaging way than is a stranger. [...] This need not be a matter of flatly denying the obvious. It is rather a matter of extending more interpretive charity to your friends than you naturally would to strangers—of offering your friends more leeway. The result, though, is that your beliefs about your friends will be slanted in their favor in various respects. (2006, 507)

Suspending judgements about the brute facts about friend's behaviour isn't the only way to protect them. It's also possible to protect a friend by suspending judgement about the moral status of their actions.¹¹ An interesting example of how our evaluative standards shift in order to accommodate our friends' misconduct can be found in *The Third Man*. At one point in the film, Anna Schmidt, Lime's lover and a woman in a good position to know about Lime's activities, tells Martins that Lime was involved in the black market. At first, Martins will not entertain this idea at all. His friend, he thinks, is no criminal! Indeed, he only accepts this when it's explained to him by Schmidt that, at least, for the Viennese, working in the black market wasn't deemed so morally objectionable as before the war. This was due to the terrible shortages in Vienna at the time. With this extra information, Martins is able to adjust his evaluative standards in a way that accommodates his esteem for his friend, Lime. Martins is thereby able to diminish his estimation of the badness of Lime's actions. It's important to note the probable reason why Martin lowered his evaluative standards. He didn't do it because he sought the truth about Lime's character, but rather, because it allowed him to protect his good opinion of Lime.

Returning to Stroud's Sam example, perhaps Sam's friends would adopt a similar strategy to deal with the accusation against Sam. According to the accusation, Sam knowingly broke someone's heart by failing to return their calls after a one-night stand. Even if his friends can't deny that Sam ultimately caused this person's heartbreak, you might suspend judgement about his blameworthiness for it. Sometimes we attempt to rationalise our epistemically partial judgements about our friends. Stroud continues:

As a good friend, you will tend to file this incident—and Sam's behavior generally—under other labels. Your reaction to this story might be, for instance, "There's never any artifice with Sam.

¹¹ On a closely related point, Ward E. Jones (2012) discusses the various strategies that are sometimes employed in order to avoid what he calls 'lover's shame,' that is, felt-shame that is induced not by your own wrongdoing, but by the wrongdoing of someone you love.

You know where you stand with him: if he doesn't want to see you, he makes that clear. There's no false politeness, no pussyfooting, no hypocrisy, no stringing you along—Sam's too genuine for any of that." In a similar vein, what other people might classify as compulsive womanizing on Sam's part, you might see as irrepressible but fickle enthusiasm and appetite for female charm in all its many varieties. (Stroud 2006, 508.)

You might even suspend your evaluative standards in order to protect Sam's well-being. Suppose that you think it's wrong to ignore someone's calls after a one-night stand. You happily apply this evaluative standard when thinking about the behaviour of colleagues you don't like and the characters from your favourite soap operas. But when you're evaluating a friend's behaviour, things are different. Sam's friends learn what Sam did, they respond by softening their evaluative standards according to which his actions are morally objectionable. We can imagine the sorts of explanations that Sam's friends might grasp for if they're forced to rationalise this move: 'Look, I know it was a bit mean of Sam not to call her back but it's not the end of the world. I mean, what did she expect? You can't expect a serious relationship to develop out of a one-night stand.' As Stroud quips, we become 'spin doctors' for our friends (2006, 508). This enables us to avoid coming to unfavourable conclusions about our friend's character.

We don't just change our evaluative standards to prevent ourselves from thinking poorly of our friends' character and conduct. There's a positive side to this too. Sometimes, we can change our evaluative standards in order to help build a positive interpretation of them. We move the goalposts, so to speak, so that our friends come out on top every time. Here's an example of how, called Drop Out: Suppose that according to Ben's prior evaluative standard, it's praiseworthy to write a PhD thesis in philosophy. He was always impressed by the fact his friend, Alana, was writing a PhD thesis in philosophy. But one day, Alana drops out of her PhD program. In order to accommodate his new belief that Alana dropped out and preserve his good opinion her, Ben suspends judgement on his evaluative standard that doing a PhD is praiseworthy. Rather, he decides, the *truly* courageous thing to do is to leave the comfort of the academic world and get a 'real job'.

Epistemic partiality of this sort seems to be a psychological possibility. We're psychologically capable of withholding belief, at least in some cases. Schroeder (2012a; 2018a) and Jacob Ross (Ross and Schroeder 2014) observe that at least when the costs of erroneously believing that p are high, we seem to exhibit a psychological capacity to withhold from believing

that p even when we possess evidence which would lead a disinterested third party to believe that p .¹² Consider the following example, from Schroeder:

[P]eople with known peanut allergies have no trouble holding out for more evidence that a sandwich is made with almond butter before concluding that this is so, than people who lack peanut allergies. Yet the difference between knowing you have a peanut allergy and not being aware of the allergy is a difference in what you are aware of about the practical consequences of believing that something does not contain peanuts. This shows that we have no trouble being sensitive to practical considerations, either, when we form beliefs – when these considerations are epistemic reasons *against* belief. (2018a, 119)

So, when the stakes are high because the costs of falsely believing that p are high, we demand more evidence before we believe that p . In other words, we set a higher epistemic standard when we process our evidence.¹³

This seems to be analogous to what we do when we cut our friends some slack, as in the cases described. When the costs of incorrectly believing something harmful about a friend are high enough, we may withhold from belief, even though we would assent to belief on the same evidence if it were not about a friend. Now, that can occur independently at the level of ‘brute facts’ (Did Sam break so-and-so’s heart?) or at the level of evaluation (Is Sam blameworthy for breaking so-and-so’s heart?). Of course, the assumption at the root of this thought is that the fact that our friendship is at risk raises the stakes of our believing incorrectly. Schroeder identifies friendships as ‘stable high-stakes cases’ (2018a). If falsely believing something about a

¹² Schroeder defends this thesis in greater detail elsewhere (see Schroeder 2012b; 2013).

¹³ It’s hard to pin down exactly what the term ‘epistemic standards’ refers to. On epistemic standards, Miriam Schoenfield writes:

Some people think of them as rules of the form “Given E, believe p !” Others think of them as beliefs about the correct way to form other beliefs. If you are a Bayesian, you can think of an agent’s standards as her prior and conditional probability functions. [...] we can just think of a set of standards as a function from bodies of evidence to doxastic states which the agent takes to be *truth conducive*. Roughly, this means that the agent has high confidence that forming opinions using her standards will result in her having high confidence in truths and low confidence in falsehoods. (2013, 199)

friend could harm them, and you value their well-being, then the stakes are raised for you. Since you don't value the relationships you have with strangers, to the extent that you even have a relationship with them, the stakes are considerably lower.¹⁴

In summary, it seems we treat our friends with epistemic partiality in the way we process our evidence about them. Sometimes we demand more evidence before we form beliefs that would harm our friends. This seems to be present in the Oluo case and in Stroud's Sam example (2006). When the evidence is mounting against a friend's character, it seems we suspend judgement rather than risk forming a belief that harms them. This can involve suspending judgements about the brute facts regarding our friends, as in the *Third Man* example, or it can involve suspending or adjusting our evaluative standards, as in Sam's Stroud example and the Drop Out example.

Even so, epistemic partiality in friendship isn't unlimited. I've already suggested a couple of cases in which it seems that epistemic partiality is morally objectionable. In the section that follows, I'll consider a number of cases in which it seems we would not treat friends with epistemic partiality.

§4 The limits of epistemic partiality in friendship

In the first part of Bernhard Schlink's *The Reader* (1997),¹⁵ Michael Berg is a fifteen-year-old boy who has an intense, sexual relationship with Hanna Schmitz, a tram guard who is twenty-one years his senior. Hanna has Michael read books to her. We later learn that this is because she's illiterate. The novel is set in Germany in 1958 as the German people struggle to come to terms with their country's recent history and in particular the atrocities of the Holocaust. A few months into their relationship, Hanna suddenly disappears. In part two of *The Reader*, Michael is a young adult, studying law at university. As part of his studies, he attends the trials of several women who're accused of overseeing the deaths of 300 Jewish women. The accused are said to have allowed the Jewish women to burn to death in a church where they were being held captive. Michael recognises one of the women. It's Hanna.

¹⁴ Schroeder (2018a) thinks that withholding belief in such circumstances is epistemically rational. In Chapter Three, where I'll discuss an epistemological objections to epistemic partiality in friendship, I'll explain why I think that this is not epistemically rational.

¹⁵ First published in German as *Der Vorleser* (1995).

Michael is horrified by what he has learnt about Hanna's past.¹⁶ The evidence is incontrovertible: Hanna, along with the other women on trial, is responsible for the deaths of the Jewish women. Michael does not give her the benefit of the doubt or look the other way. The only way to respond to what he has learnt about Hanna is accept it, like a punch in the gut. Michael struggles with this for the rest of his life.

This case gives us a lot to think about regarding the limits of epistemic partiality. Why doesn't Michael treat Hanna with epistemic partiality? There are three plausible (and indeed compatible) explanations of this. First, I think it's uncontroversial that for moral reasons, Michael shouldn't treat Hanna with epistemic partiality. Hanna's crimes are indescribably evil, there's no doubting it. No matter how much he cares for her, he has an overriding moral duty to confront the evils of the Holocaust head-on. To excuse Hanna from blame for her part in the Holocaust is deeply, morally objectionable. To put it another way, whatever moral reasons he might have for treating Hanna with epistemic partiality are outweighed by the moral reasons he possesses not to treat her with such partiality. This is a plausible interpretation of Michael's motivations. We can read *The Reader* as a story about how the post-war generation (i.e., Michael's generation) grapples with the sins of the previous generation (i.e., Hanna's generation).

Second, it's not clear Michael actually possesses moral reasons to treat Hanna with epistemic partiality. She really is morally responsible for the deaths of the Jewish women she was holding captive and she admits it. This is important. Given that she admits her guilt for the deaths, Michael's coming to believe that she's guilty is not obviously a harm. It's possible, at least, that when she admits her guilt she actually invites others to judge her accurately, that is, without undue concern for her interests. The thought is that she doesn't benefit from his treating her with epistemic partiality in this case. And if Hanna doesn't stand to benefit from Michael's treating her with epistemic partiality, then it's not clear that he has any reasons—moral or otherwise—to do so.

Third, it seems as if it's psychologically impossible for Michael to treat Hanna with epistemic partiality. There's no psychologically feasible way for him to process the evidence in his possession in a way that successfully preserves his favourable judgement about Hanna's character. For one thing, the evidence in Michael's possession after he attends Hanna's trial is, to echo Stroud (2006), incontrovertible. The evidence presented in the trial is overwhelming and

¹⁶ The case is complicated somewhat when the other women try to single out Hanna for blame. The problem is that Hanna, ashamed of her illiteracy, signed a witness statement that she didn't understand.

besides, Hanna admits her involvement in the deaths of the Jewish women. There's no denying that Hanna was responsible for the deaths of the Jewish women. When faced with overwhelming evidence of her involvement in the deaths, it seems withholding judgement is no longer a psychological possibility.

What's more, there's no way he can suspend or adjust his evaluative, moral standards about the depravity of her crimes. Though it might be possible to suspend some of his evaluative standards, this is plausibly much more difficult when those standards represent firmly-held convictions. So, in Michael's case, his moral conviction about the evils of the Holocaust is held too firmly to be easily suspended just because it would protect his view of Hanna.

We find a further illustration of the moral limits of epistemic partiality in friendship in *The Third Man*. Martins' efforts to protect his friend are cut short when he learns about the moral gravity of the accusations against Lime. Let's revisit this briefly. When Martins is under the impression that Lime stands accused of smuggling, he seeks (successfully, from his perspective) a way to excuse Lime from moral blame. When Schmidt tells him that smuggling is not considered morally wrong in Vienna at that time, he willingly accepts her testimony. For the meantime, Lime is excused from blame, at least in his eyes. This all changes after Calloway shows Martins around the hospital ward where many of Lime's victims lie dying. At this point in the film, Martins stops trying to search for evidence which would justify his rejection of the accusations against Lime. It seems as if whatever reasons he previously had for protecting his opinion of Lime are now overridden by his moral reasons not to do so.

Perhaps Martins still cares about Lime. After all, it's conceivable that he can't simply shrug off his concern for his friend's well-being. But, the viewer is given to understand, Martins is a good man at heart. As such, he can't allow himself to prioritise his friend's well-being over the well-being of Lime's victims. This suggests that there are moral limits to the amount of epistemic partiality with which we can treat people. Once he understands what's morally at stake, Martins can no longer bring himself to protect Lime. Now let's consider another example in which we find the moral limits of epistemic partiality.

Not all limits on epistemic partiality in friendship are imposed by the moral-wrongness of a particular case of epistemic partiality. Sometimes, epistemic partiality is simply unmotivated and, what's more, sometimes we have moral reasons to be epistemically impartial towards our friends. In other words, sometimes, if we really care about our friends, we will be *more* attentive to the facts about them than we would if they weren't our friends, other things being equal. We can question the claim that we should treat a friend with epistemic partiality from both a first-person perspective, where we consider what we want from good friends, and a third-person

perspective, where we consider how someone's actions reflect on them as a friend. According to Arpaly and Brinkeroff (2018, 43) from neither perspective is it clear that we should always treat friends with unlimited epistemic partiality.¹⁷ They write:

Consider Meredith and her friend Ying. Meredith is deliberating about whether or not to open up a bakery, and asks Ying for her insight on the matter. We can imagine Meredith thinking like this: "This is a big life decision, and I know Ying wants what is best for me. I want to know what Ying *really* thinks—does she honestly believe I have what it takes to run a successful small business?" (2018, 43)

The thought is that we know that epistemic partiality in friendship isn't always good for us. Sometimes we like our friends to have flattering views, sure enough, but there are times when the accuracy of their beliefs is much more important to us than whether their beliefs protect us from offence. Arpaly and Brinkeroff are correct, I think, that this is one such case.

For Arpaly and Brinkeroff, it's not only from the first-person perspective that it's unclear that friendship requires epistemic partiality. When we consider someone's conduct as a friend, it's sometimes the case that we think it's praiseworthy that someone doesn't give their friend the benefit of the doubt.

Let us take a step back and consider Meredith and Ying. Suppose it is a bad idea for Meredith to attempt to open up a bakery on her own. While Meredith bakes delicious goods, she is highly disorganized and lacks basic accounting skills—she has trouble keeping up with her own personal finances. If Ying advises Meredith not to open up a bakery on the grounds that she lacks the business savvy in order to successfully run it, and if the reason for this advice is that Ying has a realistic view of Meredith's limitations, we do not take Ying to be a lesser friend. Nor do we think that Ying would be a better friend had she instead reacted to the evidence of Meredith's limitations by taking steps to think more highly of Meredith's organizational and accounting skills. If anything, it reflects well on her character as a good friend (albeit not a good sycophant). (2018, 43-44)

This example suggests that we don't always see epistemic partiality in friendship as a praiseworthy display of loyalty and concern. Moreover, the example suggests that this can be the

¹⁷ As I'll explain later in this thesis, I think they're mistaken about how troublesome the existence of such cases is for those who defend the claim that friendship requires epistemic partiality.

case even in the absence of external moral reasons not to treat a friend with epistemic partiality, such as we saw in *The Reader* and *The Third Man*.

Finally, it's important to note that there are prudential limits to the extent to which we should treat our friends with epistemic partiality too. Sometimes, for instance, it's too risky to give our friends the benefit of the doubt. If a friend asks you to take out a £10,000 payday loan on their behalf, then you probably shouldn't give them the benefit of the doubt when figuring out whether or not to take the loan for them. After all, if it turns out that they can't pay you back, you'll be in serious financial trouble. In other situations, it would be imprudent to treat our friends with epistemic partiality, not because of the risks of error, but because of the costs of stopping, withholding our judgements, and reconsidering what we believe about them. I'll discuss such cases in more detail in the following chapter.

There are limits, then, to the extent to which we treat our friends with epistemic partiality. Some of these limits seem to be moral in character, though the sources of our moral reasons may come from within the relationship, when our friend's well-being is at stake, or without, when epistemic partiality is incompatible with some overriding moral duty. Other such limits are prudential in character. Sometimes, so the thought goes, we just can't afford to give our friends the benefit of the doubt.

§5 Summary

We treat our friends with epistemic partiality. Epistemic partiality influences the way we gather evidence about our friends and how we process that evidence. It seems that we do this, and we do it for our friends' benefit, both positive and negative. It also appears that this sort of behaviour is not limited to a few extreme cases, such as *The Third Man* case. Rather, it seems like a fairly ordinary part of our epistemic lives, as in Stroud's Sam example and Keller's Poetry Reading example. But we've also seen that there are limits to the extent to which we treat our friends with epistemic partiality. Sometimes it seems we respond to external moral reasons not to, as in the cases of *The Reader* and *The Third Man*. Other times, our care for our friends is directly the source of moral reasons not to treat them with epistemic partiality.

But this chapter has been descriptive. The goal was to establish what it means to treat our friends with epistemic partiality. Naturally, the question arises whether we should treat our friends with epistemic partiality. That, then, will be the subject of the next chapter of this dissertation.

Chapter Two

Why we should treat our friends with limited epistemic partiality

§1 Introduction

We treat our friends with limited epistemic partiality, or so I argued in the previous chapter. We saw this in Martins' treatment of Lime in *The Third Man*, when he went charging around Vienna in search of exculpatory evidence. We saw it again in Ijeoma Oluo's (2015) description of a family's struggle to accept that a family member had sexually assaulted someone, against mounting evidence. These are just two of many examples we considered. This epistemic partiality, I claimed, occurs at both the levels of evidence gathering and evidence processing.

The project in the previous chapter was descriptive. In this chapter, the project takes a normative turn. I will argue that we should treat our friends with epistemic partiality, although that partiality must be limited. My defence of this thesis consists of an argument which I call the 'argument from care'.¹⁸ According to the argument from care, we should treat our friends with limited epistemic partiality because friendship is a source of *pro tanto* reasons to do so. Beliefs, I'll argue, can harm, and epistemic partiality is the best way to protect our friends from that harm.

First, though, I'll present the theory of friendship which I'll be employing in this chapter and for the rest of this thesis. Then I'll explain why we should treat our friends with limited epistemic partiality. Finally, I'll respond to a possible objection and explain why, even though friendship requires epistemic partiality, that doesn't entail that we should treat everyone with epistemic partiality.

¹⁸ The argument from care is a moral argument for limited epistemic partiality in friendship.

Though I do not discuss it in this thesis, it's plausible that we have *pro tanto* prudential reasons to treat our friends with epistemic partiality too. The argument would go something like this: (1) If treating your friends with limited epistemic partiality promotes your friendships, then, you have *pro tanto* prudential reasons to treat your friends with limited epistemic partiality. (2) Treating your friends with limited epistemic partiality promotes your friendships. (3) Therefore, you have *pro tanto* prudential reasons to treat your friends with limited epistemic partiality.

§2.1 On friendship

If someone is a friend, how ought we to treat them? This section discusses the place of care and commitment in friendship. First, though, it's important to make some clarifications about the nature of close friendship and to mark clearly the difference between close friendship and what we might call 'instrumental' friendship. Aristotle famously distinguishes between close friendship and friendships of pleasure or utility. This distinction is no mere technicality, but rather, it reflects real differences among the kinds of relationships we call friendships. A true friend, Aristotle claims (2009, 145: 1156b10), values their friend's well-being for its own sake. Close friends find each other's well-being to be non-instrumentally valuable. In this respect, close friendship contrasts with other, solely instrumental forms of friendship. According to Aristotle (2009, 146: 1156a10-15), in friendships of pleasure or utility, a person loves their friends by virtue of either the pleasure that they bring them, as in friendships of pleasure, or else by virtue of the utility that the friendship yields for them, as in friendships of utility. In this thesis, I focus on close friendships.¹⁹ So, with this distinction in mind, how ought we to treat our close friends?

§2.2 Care and friendship

Towards the end of Toni Morrison's *Beloved*, we find Paul D comforting his friend, Sethe, after her daughter's second and final disappearance. He's reminded of these words, spoken by another friend, Sixo, about the care showed to him by a former friend:

She is a friend of my mind. She gather me, man. The pieces I am, she gather them and give them back to me in all the right order. It's good, you know, when you got a woman who is a friend of your mind. (Toni Morrison 2005, 321)

Friendship is a caring relationship. That is to say, it's constitutive of friendship that we care about our friends and their well-being. As Aristotle puts it, good friends care about each other's well-being for its own sake (2009, 144: 1156b9–11). If A is friends with B, then B's well-being is

¹⁹ My purpose here isn't to disparage instrumental friendships. We happily engage in friendships for the sake of pleasure and convenience as and when it suits us—we break off such friendships just as easily. James Grunebaum (2005) argues persuasively that philosophers have been unfairly disparaging in their assessment of the value of such friendships.

non-instrumentally valuable to A. This precludes relationships in which one person cares about the other's well-being only for instrumental reasons. Suppose B is A's best employee. If A only desires B's well-being because B makes a lot of money for A's business, then A doesn't care about B's well-being in the way required for friendship. To be clear, when a person desires their friend's wellbeing, they do so because their friend's well-being is non-instrumentally valuable to them. That's what it is to truly care about a friend's well-being.

The thought that friendship necessarily involves mutual care or goodwill has endured. Neera Badhwar (1987, 1) defines friendship as 'a practical and emotional relationship of mutual and equal goodwill, affection and pleasure'. Extending this point, Badhwar continues:

The best, most complete friendships are those in which friends love and wish each other well as ends in themselves, and not solely, or even primarily, as means to further ends--social advancement, amusement, the promotion of some cause, or even mutual edification or improvement. (1987, 1)

Elizabeth Telfer has a similarly Aristotelian view. On her account, friendship necessarily involves mutual affection. Affection, for Telfer, involves

a desire for another's welfare and happiness as a particular individual. This desire is thus to be distinguished both from sense of duty and from benevolence. (1970, 224)

It's interesting that she specifies that the desire for another's well-being is for them as a 'particular individual'. The thought is that when you have affection for someone, the source of your desire for their well-being is found in the value that you place in that very person's well-being—above and beyond what value you might also place in the value of human well-being in general.

What exactly it means to 'care' about someone's well-being is no doubt a complicated matter, involving doxastic, affective, and volitional attitudes. For purposes that will become clear, I want to focus on one aspect of caring about someone's well-being, namely, desire. To be clear, if you care about someone, then, among other things, you desire their well-being. Harry Frankfurt writes that the fact that

a person cares about or that he loves something has less to do with how things make him feel, or with his opinions about them, than with the more or less stable motivational structures that shape his preferences and that guide and limit his conduct. (1999, 129)

We might think of friendships as a kind of loving relationship.²⁰ Considering friendship as a loving relationship lends yet more plausibility to the claim that friendship necessarily involves care. Frankfurt spells this out clearly:

It is essential to all instances of the type of love with which I am concerned that: (a) the lover is devoted, in some degree non-voluntarily, to the flourishing of his beloved; and (b) he desires the well-being of his beloved for its own sake rather than only for the sake of ways in which it might support or promote other interests. (1998, 4)

The claim that care is necessary for friendship chimes well with our intuitions about friendship. We can see this more clearly by considering the reasons why a friendship might break down. Suppose that A and B are friends. Then, after some time apart, A's care for B gradually weakens until they no longer care about B at all. As A's care for B weakens, it seems like the quality of their friendship diminishes too. And when A no longer desires B's well-being, so that learning of B's hardships stirs no more sadness or anxiety in them than do the hardships of perfect strangers, it seems that their friendship with B is over.

It's plausible, then, that friendship is a caring relationship. Friends care for each other, where caring for each other entails a kind of investment in each other's well-being.

§2.3 Commitment and friendship

Good friends stand by us through thick and thin. If someone's your close friend today, then you can expect that they'll be your close friend tomorrow, next week, next month, and next year, notwithstanding when the unspeakable happens. That's because friendship requires long-term, voluntary commitment.

Specifically, being someone's friend involves a kind of voluntary commitment to care for them. The fact that friendship requires commitment allows us to explain how we can voluntarily terminate a friendship. If commitment were not necessary for friendship, then it would be hard

²⁰ The idea that friendship is a loving relationship is far from unprecedented. See, for instance, Laurence Thomas 1993. For Eleonore Stump, similarly, friendship is a distinctive 'office of love' (2010, 97-98). As such, on her view, it necessarily involves the desire for the good of beloved. Similarly, for Gabriele Taylor (1976, 153-4), 'if x loves y [...] x wants to benefit and cherish y '.

to explain how we can terminate friendships voluntarily. If you decide that you don't want to be friends with someone but your relationship remains one of mutual care and esteem (let's say), then you're still friends—whether you like it or not.

Accordingly, an explanation of commitment is necessary for a successful account of close friendship. Nevertheless, the nature of commitment in friendship is puzzling. Thomas writes that

On the one hand, there is clearly something to the idea that friendships are an expression of choice; no one supposes that she or he had no choice but to be a person's friend. Yet, it is all too obvious that as a rule we do not self-consciously choose our friends in the way that we choose, say, the clothes that we wear. One does not shop for a friend in the way that one shops for an article of clothing. There is a very clear sense in which we grow into friendships; indeed, we can even be surprised that our interaction with someone has given rise to such companion friendship. (Thomas 1987, 217-218)

While it's clear that we exercise a degree of choice when we commit to friendships, on Thomas's view, at least, it's more mysterious what exactly we choose and how we choose it. What's more, it's not obvious that there's a moment at which we choose to become friends with someone. Richard White, for instance, writes:

[I]t can hardly be denied that friendship does seem to involve a commitment to another person. This need not ever be explicitly given, or recognized at a particular time. Still, insofar as we do begin to identify the joys and sorrows of another person with our own concerns, and apparently endorse the enlargement of being by choosing continued contact with them, then we are also creating expectations about our availability to the other person, and making a commitment in the relevant sense. (1999, 81)

So what do you commit to, when you commit to friendship with someone? It can't just be a matter of choice. You can't just choose to be friends with someone because, as Telfer (1970, 230) points out, you can't voluntarily summon up the concern for someone's well-being. Rather, when you commit to friendship with someone, you commit to the caring relationship that you recognise as forming between you. To put it more formally, if A and B are friends, then they commit to making a policy of caring for each other.

For Telfer, it's a necessary condition for friendship that friends acknowledge the caring and close relationship that is starting to exist between them, and consent to it.

This acknowledgement involves, not so much the *formation* of a policy, as endorsement of or consent to a policy which is by then enshrined in practice. This is part of what is meant by *commitment* in friendship. (1970, 230)

With respect to care, this acknowledgement is a matter of recognising your concern for your friend's well-being and the activities that you undertake in order to safeguard it.

This way of thinking about commitment seems plausible. Clearly, friendships don't magically begin with the deliberate formation of and commitment to such policies of care. Rather, relationships graduate into friendships from friendly acquaintances with the acknowledgement and acceptance of policies of behaviour which are constitutive of friendship. Reflecting on the nature of commitment in friendship, White writes the following:

[W]hen I spend time with someone, accept their help, and make myself available to that person, by sharing the more intimate aspects of myself, I am also creating an expectation that is equivalent to a commitment, given the institution of friendship and all that it commonly entails. In fact, it is both a commitment to my friend, and a commitment to myself: to view her as a friend and to be open and available to her, as well as to affirm the friendship as a more or less important expression of who I take myself to be. (1999, 82)

When you commit to friendship with someone, you commit to a relationship of care and closeness with them. Making such a commitment is like making a promise, in the sense that when you make the promise, you grant your friend the right to expect you to behave in certain ways.

Friendship, then, involves care and commitment. Friends care about each other, which means, among other things, that they desire each other's well-being. One aspect of commitment in friendship is found in the friends' recognition and acceptance of what policies of mutual care develop between them.

§3.1 The argument from care: a moral argument for limited epistemic partiality

It's plausible that we treat our friends with epistemic partiality. What's less clear is whether we should. In this section, I'll argue that we should treat our friends with limited epistemic

partiality.²¹ The sense of ‘should’ involved here is moral, as will become clear soon. Accordingly, the argument from care is a moral argument.

The argument from care has the following structure:

- (1) A and B are friends.
- (2) If A and B are friends, then A cares about B.
- (3) If A cares about B, then A has a *pro tanto* moral reason to protect and promote B’s well-being.
- (4) Treating B with epistemic partiality protects and promotes B’s well-being.
- (5) If treating B with epistemic partiality protects and promotes B’s well-being, then A has a *pro tanto* moral reason to treat B with epistemic partiality.
- (6) Therefore: A has a *pro tanto* moral reason to treat B with epistemic partiality.

Friendship, I have argued, is a caring relationship. If two people, A and B, are friends, then necessarily they care about each other. Premises two and three follow straightforwardly from the account of friendship described above. When someone is your friend, you have a *pro tanto* moral reason to protect and promote their well-being. Your reason to do so is a *moral* reason because its source is located in your duty to someone else, that is, your friend. That distinguishes the kind of reasons relevant here from prudential reasons, which concern your duties to yourself. Your moral reason for protecting and promoting your friend’s well-being is a *pro tanto* moral reason because it can be over-ridden by other *pro tanto* reasons that you may possess.

The premise upon which the whole argument from care turns is premise four: Treating B with epistemic partiality protects and promotes B’s well-being.²² To defend this premise, it must first be shown that A’s beliefs can harm or benefit B’s well-being. Then it must be shown that epistemic partiality can prevent A’s beliefs from harming B’s well-being and allow A’s beliefs to

²¹ Sarah Paul and Jennifer Morton (2018) argue that it is morally permissible to be resilient to the evidence about your friends’ sincerity in their commitments.

²² It’s important to note that this explanation of how we come to possess *pro tanto* reasons to treat our friends with epistemic partiality is compatible with Piller’s view that friendship is, by and large, a historical concept (2016, 335). That is to say that the reasons we’re friends with people are mostly to do with the fact that we’ve passed lots of time together in shared activity, rather than for reasons of enduring, mutual esteem.

benefit B's well-being. Finally, it must be shown that possible alternatives to epistemic partiality are unfeasible.

§3.2 Our beliefs can harm and benefit our friends

If epistemic partiality can help us to protect our friends, then presumably our beliefs can harm them. There are, I suggest, many ways in which beliefs can harm people. What's more, our friends are especially vulnerable to such harm. However, before explaining how our beliefs can harm and benefit our friends and why they're so vulnerable, it's important to explain why it's reasonable to think that beliefs are the kinds of things that can harm or benefit anyone.

What you believe about your friends can harm or benefit them. In fact, as Keller (2004, 2018) claims, it's possible to explain how beliefs can harm or benefit on any theory of well-being. There are three main theoretical positions regarding the nature of well-being: mental state theories, desire theories, and objective list theories. Whilst there's a lot of in-camp variation, each of the three positions settles around one of the following three claims. Mental state theorists say that all that matters is pleasure and pain. Desire theorists think that your well-being depends on your getting what you want. Objective list theorists think, crucially, that your attitudes aren't relevant to your well-being—at least, not necessarily. Rather, objective list theorists claim that there are specific, identifiable goods, such that attainment of these goods contributes to our well-being. Now I'll briefly explain why it's reasonable to think that the hypothesis that beliefs can harm or benefit a friend is compatible with each of these theories.

I'll start by briefly describing a mental state theory of well-being, or hedonism, as it's sometimes known.²³ The central thought of such theories is that the only relevant facts for determining your level of well-being are the facts about the amount of pleasure and pain you have experienced.²⁴ The more pleasure you experience, the higher your level well-being is. Conversely, the more pain you experience, the lower your level of well-being is. This crude summary will suffice for now.

²³ Earlier hedonisms have been defended by Jeremy Bentham (1789/1907) and John Stuart Mill (1863/2002). More recent defenders of hedonism or mental state theory include Fred Feldman (2004), Ben Bradley (2009), Sharon Hewitt (2010).

²⁴ Mental state theorists disagree about whether quality or quantity of pleasure is more important. Quantitative hedonism is the view that what matters is the amount of pleasure you experience. Qualitative hedonism is the view that pleasures can be graded by quality.

If a mental-state theory of well-being is correct, your beliefs about someone can harm or benefit if and only if your beliefs about them can make a difference to their levels of pain or pleasure. We might cash this out more fully in terms of insult. It's plausible that your attitudes towards someone can be insulting, for example, if they disparage or show contempt towards the other person.²⁵ Let's revisit the Othello example from Chapter One, in which Othello comes to believe that Desdemona has betrayed him. Clearly, Othello's belief that Desdemona betrayed him causes her severe emotional pain. In a manner of speaking, Othello's belief about Desdemona's character is insulting to her.²⁶ In Stroud's Sam example (2006), we can imagine how Sam would be insulted if his friends accepted the allegation of cruelty against him. Likewise, in Keller's Poetry Reading example (2004), we can imagine that Rebecca would be insulted if Eric came away without a flattering view of her performance. It's also plausible that your holding certain beliefs can cause someone else pleasure. Considering the Sam example again, we can imagine that if Sam's friends ignored the third party's allegation against him, he would feel comforted by their loyalty. Similarly, we can imagine that Rebecca would be pleased to learn of Eric's flattering opinion of her poetry reading.

Perhaps the most famous criticism of mental-state theory comes from Robert Nozick, whose pleasure-machine example is taken by many to have all but sunk the prospects of mental state theory of well-being (1974, 42-45).²⁷ The experience machine example is said to show that a life spent in the machine would be devoid of many valuable features. For example, the life of the plugged-in subject would lack close, personal relationships and meaningful achievements. Surely, so the thought goes, these are very desirable features of human life. This intuition leads us to the desire theory of well-being.

²⁵ I'll address the question of what counts as an insult later in the chapter.

²⁶ Iago's deception ultimately leads to the breakdown of Othello and Desdemona's relationship. It's plausible that what you believe about someone can justify your decision to break off your relationship with them. Aristotle argues that we should break off friendships if we learn that a friend is irredeemably bad (2009, 167: 1165b10-25).

²⁷ Sharon Hewitt (2010, 331) argues that our intuitions about the experience machine are not such strong evidence for the falsity of hedonism after all, because in the real world 'seeing certain things besides pleasure as ends in themselves may best serve hedonistic ends, [so that] hedonism may justify our taking these other things to be intrinsically valuable'. Ben Bradley also defends the mental state theory from this and other objections (2009).

According to desire theories of well-being, the only facts that matter when it comes to determining your level of well-being are the facts about whether your desires have been satisfied or frustrated.²⁸ If the desire theory of well-being is correct, then what you believe can make a difference to another's well-being if your holding (or not holding) certain beliefs satisfies or frustrates their desires. That said, the simple version of desire theory faces objections. Not every satisfied desire is good for human well-being. Chris Heathwood (2006, 540-541) notes that if your desire that p is satisfied but you only desired that p because the state of affairs that p would bring about another state of affairs, q , then your level of well-being doesn't increase just because your desire that p was satisfied. From this we can conclude that what really matters are intrinsic or non-derivative desires, and not desires which are held only extrinsically or derivatively.

Moreover, people sometimes desire things which are bad for them. Heathwood (2006, 544) calls this the problem of defective desires. One solution to this problem is to exclude the satisfaction of your uninformed desires from consideration when determining your level of well-being. Only the facts about the satisfaction your informed desires matter when it comes to determining your level of well-being.²⁹ Heathwood (2006) proposes that when we consider what level of well-being someone currently has, we should include the facts about the satisfaction of their desires *all things considered*, rather than simply focussing on the one ill-informed desire.

If a desire theory of well-being is correct, then what you believe can make a difference to another's well-being if your holding (or not holding) certain beliefs satisfies or frustrates their desires. A quick survey of the examples we've discussed so far suggests that the desire theory of well-being can plausibly explain the how your beliefs can harm or benefit someone's well-being.

Desdemona desires that Othello believes she's faithful to him. When Othello is deceived, her desire is frustrated. Her anguish is evidence that Othello's belief has harmed her.³⁰ In Stroud's Sam example (2006), it's plausible that upon hearing that there are rumours circulating about his character, Sam desires that his friends don't believe that he's cruel. Moreover, we can imagine the pain he might feel if his friends frustrated this desire. In Keller's Poetry Reading example (2004), perhaps one of the reasons that Rebecca invited Eric to the poetry reading is because she desires that Eric believes she's a skilled poet. Keller (2018) explains that he desires

²⁸ Desire theories of well-being have been defended by Peter Railton (1986), Mark Murphy (1999), and Christopher Heathwood (2006, 2016). Similar views are found in Henry Sidgwick's *The Methods of Ethics* (1874) and John Rawls's *A Theory of Justice* (1971).

²⁹ John Rawls defends a version of this view in *A Theory of Justice* (1971).

³⁰ That's not to say that it's evidence in favour of desire theories of well-being.

that his running friends think he's a competent runner. Now, it's plausible that Keller's desire here is weaker than, say, Sam's desire that his friends don't think he's cruel. Interestingly, according to desire theory, the strength of your desire (in part) determines the extent to which its satisfaction contributes to your well-being and how much harm is caused by its frustration.

Finally, there are objective-list theories of well-being. Objective-list theories take a rather different tack. Roger Crisp (2006, 102) explains that objective-list theories, rather than seeking to explain what constitutes our well-being, instead seek to enumerate the factors that determine someone's level of well-being. These factors have come to be known as objective goods. According to Guy Fletcher, for instance, well-being consists in our attainment of the following goods: 'Achievement, friendship, happiness, pleasure, self-respect, virtue' (2015, 149). Paradigmatic objective-list theories of well-being tend to be pluralistic, in that they typically enumerate several goods rather than one.³¹ This again distinguishes objective list theories from mental state and desire theories of well-being.³² Some other key factors which could contribute to a person's well-being could be her health, her reputation, or her living a morally good life.³³

If an objective-list theory of well-being is correct, then what you believe can harm someone if it denies them access to some item on the list of objective goods, other things being equal.³⁴ Plausibly, other things being equal, it is objectively good to have a good reputation, and

³¹ It's worth noting that objective list theorists aren't bound to the principle of pluralism. As Shelly Kagan (1992) and Chris Heathwood (2015) point out, there's no reason why we can't even think of mental state or desire theories as objective list theories; albeit very short lists, comprising only one item each.

³² For other objective list theories, referenced in Fletcher 2015, see Finnis 1980 and Parfit 1984. It's notable that Finnis and Murphy both count friendship among the number of objective goods.

³³ A diagnostic principle, which we find in Finnis 1980, goes as follows: X is an objective good if some person's having X is better for them, all other things being equal. Consider friendship. If having flourishing friendships is better for you than not having flourishing friendships, all other things being equal, then having flourishing friendships is an objective good. Finnis (1980, 72) gives the example of knowledge: Other things being equal, he says, the informed person is always better off than the uninformed person. This rule isn't perfect, but it's a useful guide.

³⁴ Presumably a respectable list of objective goods would be accompanied by an list of objective harms. Taking inspiration from Finnis (1980), we might apply the following heuristic to identify

objectively bad to have a bad one. Likewise, other things being equal, it's better to be esteemed than not, and it's worse still to be held in contempt or mocked. Sam's reputation among his friends would suffer if they came to believe that he was cruel. If previously they esteemed his respectful nature, they may now come to hold him in contempt because of what they believe about him. Similarly, if it's better to be esteemed, then it's good for Rebecca's well-being if Eric esteems her for her poetic skill, other things being equal. It's better to be adored than not, and worse still to be resented. But when Othello comes to believe that Desdemona betrayed his trust, his love for her ends and is replaced by an overwhelming resentment. This, surely, is bad for Desdemona's well-being, and not only because it leads to her murder.³⁵

It's reasonable to conclude that your beliefs can harm or benefit another's well-being, no matter which theory of well-being is correct. Now we've answered the question of possibility, we can move on to the next question, namely, how, can our beliefs harm and benefit our friends?

§3.3 Beliefs can harm

Desdemona is deeply insulted when Othello calls her a 'cunning whore of Venice'. But it's not just what Othello says that wrongs Desdemona. Rather, his belief itself wrongs her. When he calls her a whore, he betrays his brutal and misplaced contempt for her moral character, insofar he judges it. His undue contempt for Desdemona, so expressed in his false belief, seems to be the source of this wrong. But what's harmful about beliefs like Othello's?

It's tempting to think of it as matter of valence. The thought would be that 'negative' beliefs are harmful to our friends' well-being and 'positive' beliefs are beneficial to them. The initial plausibility of this suggestion is bolstered when we consider some of examples discussed in Chapter One. On this view, Othello's belief that Desdemona cheated on him harms her because it has a negative valence. And in the Drop Out example, Ben's belief that Alana's decision to drop out from her PhD programme was bold is beneficial to her because it has a positive valence.

However, the view that the valence of a judgement is sufficient for determining whether it's harmful isn't tenable. First, not all negative beliefs and judgements are harmful. Most mature people are willing to acknowledge their own shortcomings. What's more, as Schroeder (2018a,

the items we'd expect on that list: X is objectively bad if some person's having X is worse for them, all other things being equal.

³⁵ Of course, other things are not always equal.

124) points out, sometimes your negative judgement about someone's capabilities can actually help you to reach a more charitable interpretation of their actions. Consider a variation of Keller's Poetry Reading example (2004). Suppose that Eric believes that Rebecca is a terrible public speaker, and he's right—she usually is. However, his negative judgement about her overall ability colours his impression of her when she finally gives a good performance. It's plausible that because of his negative judgement, he'll be more impressed by his friend's performance than if he started off with a positive belief. Far from harming Rebecca, Eric's negative judgement about her public speaking actually lifts her up, to borrow Schroeder's phrase.

Secondly, beliefs and judgements with a positive valence can be harmful. Schroeder (2018a, 124) illustrates how apparently positive judgements can harm with the example of Phyllis, who comes to believe that her Asian neighbour is good at maths, but only believes this because of a common stereotype about Asian people. Emmalon Davis, describing a very similar case, identifies this as a form of epistemic injustice which she dubs 'identity-prejudicial credibility-excess' (2016, 487).³⁶ While the valence of the Phyllis's belief is positive, it's nevertheless harmful to her neighbour. The valence of a belief or judgement about someone's character, conduct, or capabilities, isn't sufficient to explain why that belief is harmful or beneficial. So, the question remains, how do certain of our beliefs and judgements about our friends harm or benefit them?

For argument's sake, imagine that the rumours about Sam in Stroud's example (2006) are pure nonsense—he didn't even have a one-night stand, let alone break anyone's heart by ignoring them after the event. If his friends nevertheless accepted the third party's testimony that he cruelly broke someone's heart, then they would insult him too.

In general, it's insulting to be blamed for something for which you're not morally responsible. Our reactions to misattributions of blame can tell us something about the moral status of those attitudes. What our reactions tell us is that they're harmful.³⁷ It's appropriate to feel insulted by the misplaced criticism of others because such criticism disparages our character. Note that the misplaced, critical attitudes are what insult or disparage us. That we feel insulted by someone's attitude is just evidence that their attitude constitutes an insult. But how can our beliefs harm others?

³⁶ I'll discuss the nature of epistemic injustice in a later chapter.

³⁷ Of course, this depends on your level of self-respect. Those who've no self-respect may not *feel* insulted in these situations. Nevertheless they *are* insulted, I'll argue. I'll explore this Strawsonian insight in greater detail when I discuss the views of Basu (2019a) and Marušić and White (2018).

In the cases we've considered, it seems that if someone's belief expresses false censure of another's character or conduct, then their belief wrongs that person.³⁸ Not all censure wrongs its object. Appropriate censure identifies where we could have done better; where we failed to live up to expectations. It identifies aspects of our character and conduct which do not merit esteem, but rather, resentment or even personal rejection by others. But censure can be misplaced, for instance when we attribute wrongdoing or failure where neither is to be found. Misplaced censure can lead people to feel ashamed of us, to regret our actions, to resent us, or even to reject us, when we don't deserve it. That a subject undergoes these emotional responses indicates that they view the object of their resentment to be diminished in merited esteem.³⁹ In other words, they view the object of their resentment as being in some sense diminished in their worth. And it's as a privation of merited esteem that we should think of undue criticism as harmful. Now, it's plausible that not all misplaced criticism is equally harmful. As I'll argue below, the importance of our friends' judgements in the construction of our self-esteem suggests that we would be especially vulnerable to them.

It's worth noting that these harms can be caused accidentally, and by justified beliefs. Just because Othello has evidence for his belief that Desdemona betrayed him, that doesn't mean his judgement isn't insulting to her. The fact that Othello possesses overwhelming evidence for his false judgements would be of little consolation to Desdemona. Likewise, if Sam didn't knowingly break anyone's heart but his friends believe that he did on abundant evidence, they harm him anyway. It's possible, I suppose, that we're only morally responsible for the harm caused by our beliefs if we lack evidence for them, but that's a debate for another time.

How does Othello's belief that Desdemona betrayed him by sleeping with Cassio harm her? Plausibly, the reason Othello's belief harms Desdemona is that it expresses false moral criticism of her. Desdemona describes her distress to Emilia: 'I cannot say 'whore'. It does abhor me now I speak the word; To do the act that might the addition earn Not the world's mass of vanity could make me' (Act IV, scene ii. 166-171). To judge Desdemona as a whore is to censure her character and conduct from a moral perspective. This account of harmful belief can, it seems, explain how Othello's belief harms Desdemona. She remained loyal to Othello, and his censure of her moral character are false. So, when he comes to hold them, he harms her. Her

³⁸ Finlay Malcolm (2018a) proposes a similar account of insult, where insult is caused by disparaging attitudes.

³⁹ See, for example, J. David Velleman 2001, Michelle Mason 2010 and Ward E. Jones 2012 on the causes of shame.

reaction when she learns what Othello believes about her is telling. Othello's judgement pays her a grave insult indeed.

We are reluctant to censure people undeservedly. This can explain the motivation of people who treat their friends with epistemic partiality even when they shouldn't. One such example can be found in *The Third Man*. If Martins came to believe that Lime was morally responsible for countless, painful deaths, then he would not insult Lime. That's because Lime really is responsible for those deaths. Nevertheless, it's plausible that the risk of insulting Lime is what motivates Martins to seek evidence of Lime's innocence and to set very high epistemic standards for believing that Lime is guilty of some crime.

The view presented here seems to explain how our beliefs can harm, at least in some cases. But there are other ways our beliefs can harm, which I'll explore in the sections that follow.

§3.4 An epistemology of harmful beliefs

It's not just what Othello believes, but also how he comes to believe it that harms Desdemona. How so? Well, it seems plausible that, given their relationship, Desdemona is entitled to expect Othello to treat her with a special regard for her well-being, including her self-esteem and her self-respect. When someone's character and conduct is under moral scrutiny from their close friends, they don't expect them to treat them the same way they would a stranger. This thought recalls Stroud's claim that we 'owe [our] friends something other than an impartial and disinterested review of the evidence where they are concerned' (2006, 504). What are we to make of this?

According to Marušić and White, '[O]ne wrongs another person when one's beliefs and judgments fall short of the regard the other is entitled to expect from one' (2018, 101). Their discussion sets off from Strawson's observation (1962) that when you and another person share an inter-personal relationship, you are entitled to expect the other to treat you with a certain degree of goodwill. For instance, as Marušić argues elsewhere (2017), when you tell something to someone with whom you share an inter-personal relationship, you expect them to trust you. But what if they don't trust you? Well, you're likely to feel hurt. And these feelings aren't unjustified or irrelevant. Actually, they can be evidence that you've been wronged. These feelings, our so-called 'reactive attitudes,' often reflect our disappointed expectations of them, and in particular, of how they value us as persons—as people with interests and concerns. They write:

When we expect others to treat us in certain ways, this reflects a more fundamental expectation, or demand, that they take us—our interests and concerns—into account in certain ways. We resent it, for instance, when a person treats us in a way that shows he does not view the harm it will cause us as providing much of a reason for forbearance. (Marušić and White 2018, 101)

You don't expect others to evaluate the reliability of your testimony in what we'll call 'objective' terms. Turning to such considerations when you evaluate testimony can seem inappropriate in close, personal relationships. As this reading of Stroud's Sam example (2006) suggests, when you tell something to a friend, you expect to be believed in some sense *because of your friendship*.

To explain this, Marušić and White turn to Peter Strawson's distinction between objective and participant attitudes (1962, 9–10). These are attitudes we can adopt when we're thinking about other people's character and conduct. To take the participant attitude towards someone is to accept that your relationship consists, in part, in some reciprocal demands for goodwill and regard. To take the objective stance is to reject that aspect of the participant attitude, and instead, to accept a conception of your relationship as one that takes the other person as something to be measured or about which you can make predictions and theories. To take the objective attitude towards someone is to be prepared to consider what explains their behaviour, and what predictions you can make about them based on what you know about human behaviour. When you sincerely enter into inter-personal relations with someone, so the thought goes (Strawson 1962; Marušić and White 2018), you accept a responsibility to take a participant attitude towards them. This responsibility is a source of moral reasons.⁴⁰

So, how does this help Marušić and White construct their epistemology of harmful beliefs? The point is that when we're involved in an inter-personal relationship with someone, we expect a certain degree of goodwill from them. Their actions can express goodwill. For instance, simply holding the elevator door to allow someone else to enter expresses (or appears to express) a degree of goodwill towards them. But the goodwill that we expect from other people in inter-personal relations can extend beyond their actions, to include the beliefs and judgements that they hold. When you make judgements about your friend's conduct or character, so the thought goes, your judgements should express some degree of goodwill towards them. Indeed, your failure to do so would explain how your judgement harms them. This insight may help to explain the harm caused in at least some of the cases we've discussed so far.

⁴⁰ These must be *pro tanto* moral reasons. Secret Service operatives aren't morally obligated to treat suspected terrorists from the participant attitude.

In Stroud's Sam example (2006), we can imagine that Sam has reasonable expectations about how his friends treat him, including when they form beliefs about him. Imagine how Sam would feel if his friends accepted the testimony of the third party and came to believe that he, Sam, knowingly and without regrets broke someone's heart after a one-night stand. It's plausible that he'd feel betrayed or insulted. How did his friends betray him? Well, so the thought goes, they should have asked him before they leapt to conclusions about his conduct. Sam's feelings aren't irrelevant. In fact, according to Marušić and White's epistemology of harmful beliefs, they're evidence that Sam has indeed been wronged.

When Sam promises that he didn't knowingly break the woman's heart, his friends could take the objective stance, and considering his track record as a source of testimony, or they could look at his testimony from a participant attitude, and look for considerations favourable to him. In other words, Sam would expect his friends to take the participant attitude towards him. If this is correct, then we have another way to understand how Sam's friends would harm him if they didn't give him the benefit of the doubt. The problem is that they failed to treat him with the goodwill that he deserves given the interpersonal nature of their relationship.

This also explains how Randall's reaction to Hope's testimony harms her. Upon receiving Hope's testimony, according to which he has a severe snoring problem, Randall adopts an objective stance towards Hope. He takes his wife's testimony as something to be evaluated rather than accepted. The thing is, according to Marušić and White's epistemology (2018), it's appropriate that Hope expects Randall to adopt the participant stance toward her. This should shape how he responds to her testimony. Unfortunately for Hope, when Randall adopts an objective stance towards her and attempts to explain away her reported belief he wrongs her.

Marušić and White's epistemology (2018) seems to explain how someone's beliefs can harm their friends in a variety of the cases we've discussed so far in this thesis. But theirs isn't the only plausible hypothesis about how beliefs harm. Sometimes, our beliefs and judgements harm others because we fail to recognise their practical identity appropriately.

§3.5 Belief and identity

Whether Eric would harm Rebecca by forming a critical view of her poetic skill in Keller's Poetry Reading example (2004) seems to depend in part on how much poetry actually matters to her. Suppose she doesn't care about poetry and only signed up to the reading out of curiosity. In these circumstances, if Eric came to think that her performance was dreadful, then it's not obvious that he's harmed her. Indeed, they might laugh about the fluffed lines and awkward

silences afterwards. However, if it were an important part of Rebecca's understanding of what makes her *her*, what Christine Korsgaard (1996) calls her practical identity, then perhaps we would expect Eric to react differently. If her identity as a poet really matters to Rebecca, then we might expect Eric to show more reluctance to criticise her.

Basu (2019a), again following Strawson (1962), reflects on our expectations of other people. Among these expectations is the expectation that, when we relate to someone personally, we expect them to see us more or less as we see ourselves (2019a, 922-298). She writes that

the wish to be related to as we are, as we see ourselves, not as we are expected to be on the basis of our race, gender, sexual orientation, class, etc. is not only common-place, this common-place and common-sense thought integrates well with both our pre-theoretic and theoretical intuitions. (2019a, 928)

Consider the following: If you approach a fellow academic at a conference, you expect to be seen (and treated) as an academic. If, say, for reasons of racial or gendered prejudice, you're treated by the fellow academic as if you're a member of the waiting staff or the non-academic wife of some other conference delegate, then it's plausible that you're wronged in some way by the person who thinks this about you.⁴¹ Of course, there's nothing wrong with being a member of the waiting staff or the wife of an academic. The problem, rather, is that *that's not you*. But there's still more to explain here. Why is it harmful if people fail to see you as you see yourself?

The reason, as Basu (2019a) argues, is that we're entitled to expect people to treat us not as objects but as persons. When they don't live up to this expectation, we're wronged. When the fellow academic leaps to a conclusion about your reason for being at the conference, they do so because they didn't see you as a person, but as an object about whose properties predictions can be made, given various 'theories' about other, similar objects. But that's inappropriate—that's treating you like an object, which you're not. Note that it's inappropriate even if the statistical evidence supports the academic's inference and, what's more, even if their inference is correct. Even if you really are a member of the waiting staff or the wife of an academic *and* the statistical evidence supports the academic's inference, he still wrongs you by making that inference and coming to the correct judgement. This, remarkably, is how an epistemically justified, true belief can harm.

⁴¹ Similar examples are discussed by Tamar Gendler (2011) in her essay, 'On the epistemic costs of implicit bias'.

Basu argues that our beliefs harm other people if they're not formed from a moral standpoint. The moral standpoint is the perspective from which other people are viewed as persons rather than as objects (2019a). What is it to view others as people rather than as objects? Plausibly, this way of interacting with people is often an ordinary part of inter-personal relations. Basu (2019a, 923) notes Rae Langton's observation that

[w]e don't simply observe people as we might observe planets, we don't simply treat them as things to be sought out when they can be of use to us, and avoid when they are a nuisance. We are, as Strawson says, involved. (Langton 1992, 486)

Naturally there's much more to be said about this observation. Crucially, though, is the thought that in your dealings with other people, you shouldn't consider them only as possible means to ends or obstacles to overcome, relative only to your own interests. When you're involved with another person, Basu writes, echoing Langton (1992) and Strawson (1962), you should accord 'a certain importance to [their] attitudes and intentions towards ourselves and [we are] cognizant of those demands with regard to our treatment of [them]' (2019a, 923).

Now, Basu is primarily concerned with explaining the harm of racist beliefs and her account offers a plausible explanation of that phenomenon. In cases of racial prejudice, the victim is viewed not as a person with her own intentions and attitudes but rather as the member of a racial group, which, it is supposed by the racially prejudiced thinker, has predictable traits and behaviours. Basu's account of harmful belief could also explain how our beliefs could harm our friends. Consider Keller's Poetry Reading example (2004). We can imagine that the notion that she's a poet is central to Rebecca's practical identity. Plausibly, part of this aspect of her practical identity is the notion that she's in praiseworthy or at least minimally competent in her capacity as a poet. So, when Eric goes to watch Rebecca at a poetry reading, he should recognise and value the fact that Rebecca sees herself in this way. When he forms judgements about her qualities as a poet, as he inevitably will, he should be responsive to what he recognises as her interests and her wishes. In this case, what's relevant is her wish to be recognised as she sees herself, that is, as a poet. Equipped with Basu's account of harmful beliefs, we're now in a position to understand why, if Eric fails to see Rebeca in a way that at least partially reflects the way she sees herself, he harms her.

Basu's account of harmful beliefs may help us to explain at least some cases in which our beliefs can harm our friends.

§3.6 On the special vulnerability of friends

In *Othello*, Desdemona is harmed by Othello's belief that she betrayed him. Plausibly, she's much more vulnerable to harm from that belief because it's Othello who holds it. Perhaps it would hurt her feelings, if some stranger fell for Iago's deception. But the fact that Othello is not just anyone, but rather, her husband, seems to explain why his belief is so damaging. Similarly, in Keller's Poetry Reading example (2004), we can imagine that Eric's judgement matters much more to her than the judgement of a stranger. Sure, she'd be a little embarrassed if some non-friend in the audience came to disparaging conclusions about her performance. But Eric's opinion, we might think, means so much more than that of the non-friend, and the reason, plausibly, is that Eric is her friend. In short, it seems that people are especially vulnerable to harm from the beliefs of their friends.

There are two reasons for this. The first reason is that friendship is a uniquely valuable source of self-esteem. The judgements and beliefs of friends seem to matter much more to us than the judgements of non-friends, other things being equal. It's plausible that a friend's disparaging judgement about your character or conduct harms you badly, even though you wouldn't be harmed to nearly the same degree if a non-friend came to hold the same judgement. Imagine you've just played and lost a game of tennis doubles. Your partner is a friend but this is your first game together. Your opponents aren't friends. Would it harm you, if your opponents came away from the game with an unreasonably disparaging view of your ability as a tennis player? I think not. In fact, in this context, you might even expect them to have an unreasonably disparaging view of you and your ability. But what if your friend and teammate had an unreasonably disparaging view of your performance in the game? That might hurt. It seems as if it's much more harmful for you that a friend would have an unreasonably disparaging view of your abilities than if, say, a non-friend had the same view.

Why can the judgements and beliefs of our friends be so much more harmful to us than the judgements and beliefs of non-friends? The disparaging judgements of our friends hurt us more than those of non-friends because of the special importance of our friends' judgements to the formation of our self-esteem. It's reasonable to assume that having high self-esteem is good and having low self-esteem is bad. Other things being equal, your level of well-being increases as

your self-esteem increases, and it decreases as your self-esteem decreases.⁴² But let's bracket those concerns for now and return to the problem at hand.

Plausibly, friendship makes a special contribution to our self-esteem. How so? As David Annis writes:

The core elements of friendship—liking, sharing, altruistic caring, and trust—all foster self-esteem. Knowing that another person likes us, wants to be with us, is concerned about our welfare for our sake, and deems us trustworthy promotes the belief that we are a person of worth and importance. (1987, 350-351)

Your self-esteem is contingent on what valuable qualities you take yourself to have. We can learn a lot about what qualities we have (and what qualities we lack) by reflecting on our friendships.⁴³ The fact that your friend likes you is evidence that you are likable, that is, that you possess properties that are worthy of esteem. Likewise, the fact that they share with you, care for you, and trust you, is evidence that you are a person worthy of sharing, caring, and trusting. Equipped with such evidence about your good qualities, so the thought goes, you're in a better position to esteem yourself. Thus, friendship contributes to your self-esteem.

Possibly, which particular qualities matter most here are determined by your practical identity, that is, the description under which you value yourself (Korsgaard 1996). If you value yourself as a trustworthy person but not as, say, an athletic person, then how positively you judge your trustworthiness contributes more to your self-esteem than how you positively you judge your athleticism. Your friends' judgements make a special contribution to your well-being. We can explain this in straightforward epistemic terms. Given that (a) your friends know you better than most strangers, and (b) your friends are people whose judgements you hold in high esteem, the fact that they esteem you would seem to be very good evidence indeed that you ought to esteem yourself.

⁴² Of course, further constraints must be made to shore up this assumption. It's highly plausible, I would suggest, that having unmerited (i.e., epistemically unjustified) self-esteem doesn't contribute to your well-being.

⁴³ John M. Cooper (1977) argues that Aristotle, in his discussion of friendship in the *Nicomachean Ethics*, sees friendship as a uniquely valuable source of self-knowledge. It's plausible that self-esteem is a form of self-knowledge and one which, moreover, friendship is uniquely well-positioned to grant us.

The judgements of our friends are special in their capacity to damage our self-esteem and, thereby, diminish our well-being. Crucially, the disparaging judgements of our friends are more damaging to our self-esteem than the disparaging judgements of non-friends. This is simply because of the evidential role played by your friend's judgements in the justification of your self-esteem is greater. Your friend knows you well and whose judgement you hold in high regard. Deprived of this evidence, you would lose much of the evidential support you have for your self-esteem. Furthermore, when you learn that a friend (as opposed to a non-friend) holds a disparaging judgement of your character, you gain particularly strong evidence that you, in fact, do not merit esteem. It's hard to maintain self-esteem when you doubt that you merit esteem from others.

To summarise: your beliefs can be more harmful to your friends than to non-friends, in part because your judgements can seriously damage the self-esteem of a friend, but not a non-friend. In Keller's Poetry Reading example, for instance, Rebecca would be harmed more by Eric's disparaging opinion of her poetry than by a stranger's. Why? In part, because, as Rebecca's friend, Eric's good opinion is much more important to her self-esteem than the opinion of a non-friend. For this reason, while the disparaging opinion of a non-friend may constitute a slight or a mild insult, the disparaging opinion of a friend is much more harmful. When our holding particular beliefs and judgements about someone's character or conduct doesn't imply a risk of harm, it's not the case that we should treat them with epistemic partiality.

§3.7 Breaking commitments and compounding harms

Now, it's plausible that failing to treat you friends with epistemic partiality can compound the harm that your disparaging judgements of their character may cause them, because your failure to treat them with epistemic partiality expresses a harmful disloyalty towards them. You owe this loyalty to your friends but, in the absence of other special relationships, not to non-friends. Loyalty can provide additional, *pro tanto* moral reasons that we should treat our friends with epistemic partiality. But, since we don't owe non-friends loyalty, loyalty cannot provide a reason for us to treat them with epistemic partiality.

Let's continue with Stroud's Sam example (2006). Let's suppose that Sam's friends don't give him the benefit of the doubt and, as a result, come to form harmful judgements about his character. Well, it's plausible that when Sam learns what his friends think about his character, he will feel hurt, not only by what they believe about him, but also by their failure to give him the benefit of the doubt. Indeed, he might feel the pang of betrayal even if his friends' judgement is

correct. This because, even though it's true that he cruelly broke someone's heart, his friends still failed to give him the benefit of the doubt, which, he might think, they owe him. We can imagine Sam's complaints: 'Look, I know what I did was wrong, but couldn't you have waited to hear my side of the story? We're supposed to be friends. I thought you'd give me a chance.' In the paragraphs that follow, I'll attempt to make sense of Sam's complaints.

First, I think this story is realistic. It's plausible that Sam would expect his friends to give him the benefit of the doubt, and that he wouldn't expect the same treatment from non-friends. The reason for this is that friendship requires loyalty and, by failing to give him the benefit of the doubt, Sam's friends have acted disloyally. It's Stroud (2006) who first connects epistemic partiality in friendship with considerations of loyalty. She writes:

It is commonly accepted that loyalty is one of the constitutive elements of friendship. One particular aspect of that loyalty, it seems, is sticking up for your friend. If someone pokes fun at one of your friends, or makes a joke at his expense, or says something false or misleading (and derogatory) about him, it seems to be your responsibility as a good friend not to join in the fun; not even simply to remain silent and withhold comment; but, on the contrary, to stick up for your friend. One might say you owe it to your friend to defend him and his reputation in the court of public opinion. The behavior of someone who simply sits silent and uneasy while others are indulging in cruel gossip about Miss X is not that of a true friend of Miss X: it is *disloyal* to join in the joke, or to let the false accusation pass unchallenged. There thus seems to be a demand of friendship to stick up for your friend. (2006, 503)

Crucially, someone's loyalty to their friend seems to involve a commitment to protect and promote her friend's interests, in spite of the costs.⁴⁴ Stroud claims that loyalty is regarded as a constitutive element of friendship. This is coherent with the account of friendship proposed at the beginning of this chapter, according to which, friendship involves making a commitment to the core conditions of a friendship, namely, care and closeness. Loyalty, I suggest, is closely related to these commitments. Let's revisit the notion of commitment to care in friendship: If A and B are friends, then A and B each recognise the caring relationship that is forming between them and commit to caring for each other. Here's where loyalty enters the picture. If loyalty

⁴⁴ Kleinig writes that 'loyalty is centrally constituted by perseverance in the conditions undergirding our relationship to an object that has come to be valued for its own sake' (2014, 32). The goal of loyalty, Kleinig claims (2014, 18), is 'securing or at least refusing to jeopardize the interests of its [i.e., loyalty's] object as part of expressing and/or further an association'.

requires commitment to another's well-being, then we can say that A is loyal to her friend, B, only if she keeps her commitment to care for B. If she reneges on her commitment to care for B without good reason, then she is disloyal to B.

The reason we should treat your friends with epistemic partiality is that it's only by treating them with epistemic partiality that we can protect them from the kinds of preventable harm described in this chapter. If A fails to treat B with epistemic partiality when she should, in fact, treat her with epistemic partiality, then A reneges on her commitment to care for B. Therefore, if A fails to treat B with epistemic partiality, A is disloyal to B.

We're now in a position to understand why Sam would expect his friends to treat him with epistemic partiality and, moreover, why it harms him if they do not. If Sam's friends come to believe that he knowingly broke someone's heart but they're mistaken, then he will suffer some harm as a result. Given that Sam's their friend, Sam's friends are committed to care for his well-being. They should take extraordinary measures to avoid falsely believing that he knowingly broke someone's heart. I have argued that the specific measures they should take are to treat him with epistemic partiality, specifically, in the way they gather their evidence, process their evidence, and update their beliefs in response to their evidence. I suggested that Sam would expect his friends to ask for his side of the story before forming a judgement about his character and conduct. This involves at least partiality in evidence gathering and belief-updating: He expects them to seek out his testimony as part of their investigation and he expects them to withhold judgement until they've got it. The reason Sam expects this partiality from them is because, as his friends, he expects them to be committed, to some degree, to protect and promote his well-being. Friendship, after all, is a caring relationship. When Sam's friends fail to treat him with epistemic partiality, they fail to keep their commitment to protect and promote his well-being. Their failure to keep this commitment is a form of disloyalty. As such, their failure harms Sam, compounding the harm that may be caused if they form harmful judgements about his character.

§4.1 Harmful beliefs and epistemic partiality

What you believe about your friends can harm your friends. What's more, your friends are especially vulnerable to such harm. But treating B with epistemic partiality protects and promotes B's well-being, or so I'll argue. To recap, epistemic partiality can occur at the levels of evidence gathering and evidence processing. Epistemic partiality at the level of evidence gathering occurs when you seek evidence that could support beliefs that benefit your friends, or

you avoid evidence that could support beliefs that would harm your friends. Epistemic partiality at the level of evidence processing occurs when you apply differential epistemic standards in order either to enable yourself to form beliefs that benefit your friends, or else to avoid forming beliefs which harm your friends.

As we saw in the case of Othello and Desdemona, forming unduly censorious judgements about another's moral character can be deeply insulting. If Othello had only withheld judgement about Desdemona's behaviour, he could have avoided coming to believe that she had betrayed him. In avoiding forming that belief, he would prevent the harm that his belief deals to Desdemona. Thus, treating Desdemona with epistemic partiality would protect her well-being.

In the *Third Man* example, Martins mitigates the risk of forming an unduly critical view of Lime's character by suspending judgement and seeking exculpatory evidence. This amounts to treating him with epistemic partiality, which prevents him from believing what ultimately turn out to be even more unpleasant facts about Lime's moral character. Now, it turns out that Lime deserved moral censure—he deserves all of our moral contempt. But Martins was reacting to a perceived risk, namely, the risk of harm by falsely forming censorious judgements about Lime's moral character. Martins treats Lime with epistemic partiality at the level of evidence gathering when he searches for exculpatory evidence regarding the allegations against Lime. He treats him with epistemic partiality at the level of evidence processing when he maintains his prior judgement about Lime's character until he finds utterly overwhelming evidence to the contrary.

In Stroud's Sam example (2006), Sam's friends will come to form harmful judgements about his character if they accept the testimony of the third party. While the strength of that testimony might be sufficient to support those judgements, they should nevertheless withhold their belief until they've heard his side of the story. Since they're his friends, he has defeasible moral justification for expecting a kind of preferential treatment from them, where his interests are prioritised, at least within limits. In cases like this, where the judgements they reach could harm Sam, they owe him more than a disinterested review of the evidence. Rather, they ought to withhold judgement until they've sought out his version of events. Again, we see epistemic partiality at the levels of evidence gathering and epistemic partiality.

In Keller's Poetry Reading example (2004), the fact that it matters to Rebecca that she be treated, in some sense, as a poet, should make a difference to how Eric forms beliefs about her. To some degree, he should seek to reinforce a view of Rebecca which concurs with her understanding of her practical identity. As we see in the example, this can involve looking for

evidence in a way that confirms the view she expects him to have of her. This, finally, is also an instance of epistemic partiality at the level of evidence gathering.

In summary, treating our friends with epistemic partiality protects and promotes their well-being. As I'll explain in the following sections, this partiality must be within limits—it's not the case that all things considered, treating friends with epistemic partiality protects their well-being.

§4.2 Pretending to believe isn't sufficient

It's tempting to suggest that, instead of treating our friends with epistemic partiality, we should just tell them white lies or otherwise pretend that we don't hold the relevant, harmful beliefs about their character and conduct. This would circumvent any need to treat our friends with epistemic partiality, or so the thought goes. So, if the possibility of telling friends white lies remains open, the argument from care is invalid. In this section, I'll briefly discuss and ultimately rebut this line of response to the argument from care.

How would this sort of strategy play out in practice? Let's consider a couple of the examples we've already discussed. In Stroud's Sam example (2006), Sam's friends receive information from a third party, who tells them that Sam knowingly broke someone's heart by ignoring their calls after a one-night stand. The thing is, this is quite a damaging allegation. If they accept it, then that might hurt Sam's feelings, especially if it's false. So, out of concern for Sam's well-being, his friends pretend to reject the third party's testimony. Truth be told, they don't know whether the allegation is true. Perhaps they suspect that it might be. But when Sam asks them if they've heard the rumour that's circulating about him, they assure him that they don't believe a word of it. Indeed, they say, they'd never believe he'd do something like that—they're his friends! And in Keller's Poetry Reading example (2004), Eric watches Rebecca while she performs her poetry at an open-mic night at a local bar. Now, let's assume that her poetry was distinctly mediocre, and Eric comes to just that judgement after watching her for a while. What should Eric do? Well, one option for Eric is to tell a white lie. He could tell Rebecca that he thinks she gave a good poetry reading even though he found it mediocre. By telling Rebecca this white lie, Eric would spare her feelings, or so the thought goes.

These strategies might seem promising. It seems that by telling white lies, we can avoid some, if not all, of the harm that could be caused by our beliefs. Unfortunately, this response is weaker than at first it might seem. First, telling white lies to spare your friends' feelings isn't always as practicable as it may seem in these examples. There's no guarantee that your efforts to

mislead your friends about what you think about them will be successful. Some people—and perhaps our friends are often an exemplary case—are very good at telling whether we’re telling the truth. In other words, you have to be a very skilled liar to protect your friend from the harm that could befall them if they found what you really think about them.

But this seems to lead us straight into another problem with the white-lie response to the argument from care. The problem is simply that lying to your friends is itself a kind of wronging. So, even if telling a white lie will help protect your friends from the painful feeling of subjective insult that they would suffer if, for instance, they learnt that you believed that they’re poetry performance was mediocre, it won’t suffice to protect them from harm altogether. To lie under these circumstances is to replace a harm of one kind of with a harm of another kind. If it’s ever to be the case that you should tell your friend a white lie, lying to them must constitute a lesser harm than that which would be caused by your belief, or so one might argue. All this makes telling white lies a high-risk strategy. The worst case scenario would seem to be the following: You try but fail to deceive your friend regarding your beliefs about their character and conduct. In this scenario, they are harmed twice—first by your beliefs, second by your (attempted) deceit. For these reasons we should be sceptical about endorsing the white lies response to the argument for limited epistemic partiality in friendship.

But even this concedes too much to the white-lie response to the argument from care, or so I’ll argue in the following paragraph. In fact, the possibility of preventing the harm of harmful belief by telling white lies is excluded by each of the hypotheses discussed in the section 3. It’s important to distinguish very clearly between the hurt of subjective or felt insult and the harm of undue censure. We’re harmed by undue censure because it disparages us. If a friend thinks you’re cruel but you’re not, then, their judgement harms you to the extent that their judgement deprives you of merited esteem. This harm is caused whether or not you actually find out about someone’s criticism. If someone tells white lies to their friend in order to disguise their undue, censorious judgements, they might be able to prevent the friend from feeling insulted, but they cannot prevent the harm caused by their beliefs.

According to Marušić and White’s account of harmful beliefs, we harm people when we fail to show them appropriate goodwill in the way we form beliefs about them. This occurs, for instance, when we fail to take them at their word or we seek evidence about their behaviour when it’s inappropriate to do so. Again, the failure to treat people in this way doesn’t harm them only if they find out about it. Rather, the source of the harm is located in the fact that when we adopt the objective stance towards a friend, we deprive them of something that they’re owed. Now, telling white lies may help us to avoid verbally insulting our friends by telling them things

they don't want to hear. But white lies can't prevent the harm that's caused by the lack of goodwill in the first place.

Finally, according to Basu's (2019) account, our beliefs harm others when we fail to see them as they see themselves. Now, if it's correct that the wrong is caused by our failure to see our friends as they see themselves, then pretending to see them in that way doesn't protect them from that wrong. If a conference delegate found out after the conference that, until they'd delivered their paper, they'd been mistaken for a member of the waiting staff, and for racist or sexist reasons, then they would likely be offended. But the hurt of subjective insult is preceded by the harm that's caused by the failure of the conference attendees to relate to the academic as they reasonably expect to be related to.

It seems that white lies can't prevent A's beliefs from harming B, even if they could help A to protect B from the hurt of subjective insult. Accordingly, epistemic partiality remains the best way to protect and promote B's well-being.

§4.3 Forgiveness and hope as alternatives to epistemic partiality in friendship

Sometimes, the mark of a true friend is not that they don't see your faults, but rather, that they forgive you. Other times, a true friend doesn't ignore your weaknesses, but hopes that you'll overcome them, and encourages you to do so. You might think, therefore, that defenders of epistemic partiality in friendship have overstated the importance of belief and underestimated the role of other, putatively non-doxastic responses to our friends' behaviour. As Kawall writes:

Among the elements of friendship are such things as acceptance, hope, encouragement, and forgiveness. With our best friends we hope and expect that they accept us, flaws and all; they see enough value in our other traits, or enough potential to change that they stand by us. They will hope that we can improve, and will presumably encourage us to do so. They will also typically be more forgiving than others when we do fall short. But all of these important aspects of friendship only come into play as we recognize our friend's shortcomings, and they ours. Indeed, it is often taken as a mark of best friends that they are quite aware of our flaws yet they still find us worthy of love; they recognize our strengths and potentials. (2013, 356-357)

The thought here seems to be the following: The cases which are said to motivate epistemic partiality in friendship don't, and the reason they do not is that we have other, non-doxastic responses to our friends at our disposal in such cases.

This response seems unpersuasive. First, as in my response to the white-lies response, it's plausible that there are many ways in which a person's beliefs can harm independently of what other attitudes they hold or actions they perform. Second, it's not clear that acknowledging the value of these non-doxastic responses encroaches on the value of epistemic partiality in friendship. Sure enough, sometimes forgiveness and hope, to take those examples, are important aspects of being a good friend. But it's not clear that forgiveness and hope are non-doxastic. Plausibly, for example, forgiving someone for wronging you requires you to judge that they no longer present a threat (Hieronymi 2001). And hope that p is often regarded as requiring the belief that it's possible that p , or at least the absence of the belief that it's impossible or highly improbable that p , at least given what the hopeful person knows (Pettit 2004). In order to show that forgiveness and hope remove or limit the extent to which we have reasons to treat our friends with epistemic partiality, it must also be shown that neither forgiveness nor hope involve doxastic attitudes. For all we know, they do involve doxastic attitudes. For that reason, the response fails.

§4.4 *Pro tanto* moral reasons for limited epistemic partiality in friendship

Premise four is correct, or so I've argued: Treating B with epistemic partiality protects and promotes B's well-being. The fifth premise of the argument goes as follows: If treating B with epistemic partiality protects and promotes B's well-being, then A has a *pro tanto* moral reason to treat B with epistemic partiality.

If premise five is true, then we can conclude that A has a *pro tanto* moral reason to treat B with epistemic partiality. There are reasons to think that is, indeed, true. Premise five follows from premises three and four. Premise three maintains that if A cares about B, then A has a *pro tanto* moral reason to protect and promote B's well-being. According to premise four, treating B with epistemic partiality protects and promotes B's well-being. Therefore, it seems to follow that A does indeed have a *pro tanto* reason to treat B with epistemic partiality.

§5.1 The limits of epistemic partiality in friendship

Treating friends with epistemic partiality doesn't always protect and promote their well-being. While treating your friend with epistemic partiality can mitigate the risks of the harms described above, it's possible that it has other, harmful consequences, which offset this benefit. Just because you would avoid harming your friend if you treated them with epistemic partiality, that

doesn't entail that it's rationally permissible to do so, all things considered: There are other ways you can harm someone besides showing them wronging them with undue censure, a lack of goodwill, or failing to treat them as a person. There must be limits to extent to which a good friend should treat her friends with epistemic partiality.

Sometimes, we have to be cruel to be kind. The truth can be the best corrective, especially in matters of personal importance. This is the first place we should look for the limits of epistemic partiality in friendship, then. Sometimes, treating someone with epistemic partiality doesn't do them any good at all. Consider the following case. Alana and Ben are at the pub. As Alana reaches for her wallet to buy another round of drinks, she notices that Ben is leaning quite heavily on the bar. Seeing Alana's concern, Ben assures her that he's not drunk too much. Alana believes him. Perhaps this is for the best: Alana would insult Ben if she didn't take his word for it. But what if Alana knew that Ben had to drive two hundred miles to London early the next morning? In these circumstances, since Alana cares for Ben, she shouldn't give him the benefit of the doubt. In fact, you might think, if she cares about Ben's well-being, she should treat him more strictly than she would a stranger, knowing that he has to make a long drive tomorrow. And note, it's directly a result of her care for Ben that she shouldn't treat him with epistemic partiality in these circumstances. Rather, she should set a very high epistemic standard, one which will probably lead her to reject his testimony. Sure enough, it might offend Ben in the heat of the moment, but all things considered, Alana has overriding moral reasons to stop him from driving regardless of whether it hurts his feelings to be accused of drunkenness or recklessness. We can think of these as inward-facing moral reasons not to treat a friend with epistemic partiality, because they relate to other moral demands generated by the friendship in question.

Here's another example. Meet Saira. Saira had always wanted to go to university but she didn't have the chance to apply when she was at college, owing to circumstances outside of her control. Now, ten years later, she's decided to apply to study literature at a very selective university. She asks her friend, Nour, who's just finishing her PhD in literature, to look over her application. Imagine that for Saira, it's now or never: If she doesn't get accepted on to her chosen course this time, she'll give up her academic dream for good. With that in mind, let's consider how Nour should treat Saira's application as she reads over it for her. On the one hand, it's plausible that if Nour comes to a critical judgement about Saira's application, she might hurt her feelings. Of course, Nour doesn't want that. On the other hand, though, this application really matters to Saira. It's been her dream to study at university and her chances of living that dream rest, in no uncertain terms, on the success of her application. In this case, it's plausible that Nour should be thorough in her evaluation of Saira's application. She shouldn't skim over

weaker sections looking for signs of strength. No, Nour should be even-handed, perhaps even strict in the way she evaluates Saira's application.

There are also outward-facing moral limits to epistemic partiality in friendship. Remember, we have *pro tanto* moral reasons to care for our friends and, therefore, to the extent that it protects their well-being, treat them with epistemic partiality. But *pro tanto* moral reasons can be over-ridden by other moral reasons. In *The Third Man*, Martins has a *pro tanto* moral reason to treat his friend, Lime, with epistemic partiality. That's because, plausibly, if he came to falsely believe that Lime was a criminal, he would wrong him. However, when Calloway finally shows Martins that Lime isn't suspected of smuggling cigarettes but rather of diluting penicillin with deadly consequences, Martins comes to possess additional, overriding moral reasons. If he still counts Lime as a friend at this point, then he still possesses *pro tanto* moral reasons to treat him with epistemic partiality. Plausibly, though, those moral reasons are over-ridden by the moral reasons he now possesses not to help Lime escape justice. The nature of the outward-facing moral limits to epistemic partiality in friendship will be discussed in Chapter Four, where I'll consider how such reasons become incumbent upon us and what limits they impose on epistemic partiality in friendship.

What this suggests is that there are cases in which care doesn't motivate epistemic partiality in friendship. While care in friendship provides *pro tanto* moral reasons which can motivate treating a friend with epistemic partiality in friendship, it can also provide moral reasons which limit it.

§5.2 Prudential limits to epistemic partiality in friendship

Sometimes, giving our friends the benefit of the doubt isn't the wisest thing we can do, given our circumstances. Indeed, sometimes, it's better for our interests if we're highly attentive to the facts about our friends' character and conduct. Accordingly, the extent to which we should treat our friends with epistemic partiality is limited. In particular, sometimes we have overriding prudential reasons not to treat our friends with epistemic partiality. In this section, I'll discuss a couple of examples to flesh out this point.

Meet Carlton. Carlton, like Saira, has always dreamt of getting a literature degree. Unfortunately for him, he has to take out a bank loan to fund his studies. Before they hand over any money, though, the bank demands that Carlton nominate a guarantor. If Carlton doesn't finish his course and doesn't have the means to pay back his loan, then his guarantor will have to pay for him. Carlton asks his good friend Alicia whether she will agree to be his guarantor. This

puts Alicia in something of a bind. In agreeing to be Carlton's guarantor, Alicia would be taking on a serious financial risk. She should consider whether Carlton can really be relied upon to complete his course and pay back his loan to the bank. The thing is that Carlton has started university twice before, and twice he's dropped out. Now he wants to try a third time. What seems to be the problem here? The thing is, if Alicia comes to believe that Carlton can't be trusted to complete his course or pay back his loan, her belief plausibly constitutes a harm to Carlton. Her judgement would be disparaging of him and, in particular, of his ability to see the serious commitment of a university course through to its end. But it's not like Alicia has nothing to lose by agreeing to be his guarantor. In fact, she stands to lose quite a lot.

Should Alicia treat Carlton with epistemic partiality in this situation? It's plausible that she shouldn't. The risk is too great. If she gives him the benefit of the doubt and agrees to be his guarantor, then she takes on a serious risk. Indeed, the risk involved in treating Carlton with epistemic partiality outweighs her interest in promoting his well-being. In other words, her own prudential reasons against treating him with epistemic partiality, that is, against cutting him some slack by applying weaker epistemic standards when considering his level of commitment, outweigh whatever *pro tanto* moral reasons she may have for treating him with said partiality. Again, we reach the limits of epistemic partiality in friendship. The case of Carlton and Alicia is an example of when, all things considered, we shouldn't avoid the ugly truth about our friends, even though it might hurt them. It's the risk of error that provides Alicia with those prudential reasons. But there are other possible sources of prudential reasons against treating a friend with epistemic partiality, in particular, the costs implicated by the very processes of epistemic partiality. Withholding judgement can make it difficult to make decisions in a timely manner. Looking for more evidence can waste even more precious time. Such delays can themselves be very costly. It's plausible, I think, that there are cases in which such considerations provide us with *pro tanto* prudential reasons against treating our friends with epistemic partiality.

Meet Claudia. Claudia has been working on her PhD thesis for four years and it's due in tomorrow. Claudia's friend, Pierre, offers to look over her thesis for her. The thing with Pierre is, he fancies himself as a bit of a scholar—he thinks that he will be able to offer some really useful feedback to Claudia. But Claudia's not so sure about Pierre's capabilities. She'd not considered it before, but her evidence about Pierre suggests that he doesn't have the great scholarly mind that he seems to think he has. Claudia starts to worry about whether she should let Pierre read her thesis. To make matters worse, she knows that Pierre really identifies as a scholar. It's a central part of how he values himself. The upshot of this is that that, if Claudia doesn't give him the benefit of the doubt, she'll probably hurt his feelings. This, possibly, is a

source of *pro tanto* moral reasons to treat him with epistemic partiality. But here's the thing: Claudia's thesis is due tomorrow, and there's no way she can afford to send it in late. In these circumstances, it's plausible that Claudia shouldn't give Pierre the benefit of the doubt. The reason why not isn't the risk of error. Actually, it doesn't matter whether Pierre could offer her good feedback. What matters is the fact that the thesis is due tomorrow. She simply doesn't have time to reconsider whether Pierre can offer her useful feedback or not.

What we see in this case is that, sometimes, we have prudential reasons not to treat our friends with epistemic partiality that come not from the possible costs of giving them benefit of the doubt and being mistaken, but rather, from the actual costs of giving them the benefit of the doubt at all.⁴⁵

§5.3 Divided loyalties: a problematic case

Before concluding this chapter, I'll briefly introduce and then respond to a more problematic case. What should you do if one of your friends tells you about the misconduct of another friend? In particular, consider a case in which it would harm the speaker if you refused their testimony, but the content of their testimony is such that if you accept it, you'll harm the other friend. In such circumstances, we seem to find ourselves with divided loyalties. On the one hand, you have *pro tanto* moral reasons to accept your friend's testimony, in order not to insult them. On the other hand, you have *pro tanto* moral reasons to refuse their testimony for the benefit of the other friend. If not a dilemma, this is at least a serious moral bind in which to find yourself.

Really, there are two variations of this case to consider. In the first variation, all three parties form one group of friends. In the second variation, though you are friends with both the speaker and the third friend, they are not friends with each other. I'll consider each variation in turn.

Should you accept your friend's testimony in the first variation of the case? Let's focus on the *pro tanto* moral reasons you have for accepting your friend's testimony, setting aside purely epistemological considerations. We might think that this isn't a true case of divided loyalty, since

⁴⁵ Of course, whether this works quite like this depends to some extent on the relationship between moral and prudential reasons. Schroeder (2018a), for example, thinks that moral reasons are simply a species of pragmatic (i.e., prudential) reasons. If that's right, then it's easy to see how prudential reasons can count against our moral reasons for treating our friends with epistemic partiality, thereby limiting the extent to which we should treat them with epistemic partiality.

you and the testifying friend share your concern for the well-being of the third friend. If you know the testifying friend and the third friend are themselves friends, then you possess epistemic reasons to believe that the reporter has the similar care for your third friend as you. When a friend tells you about a third friend's testimony, you're likely to be in possession of additional, relevant information about their motives for telling you. When a concerned friend reports the misconduct of a third friend to you, you have additional epistemic reasons to believe that they're telling you out of concern for the third friend, and not, for example, for the sake of idle gossip or worse. Therefore, if a friend tells you about the misconduct of a third friend, you should accept their testimony.

The second variation of this problematic case is more difficult to resolve. In these circumstances, it's probable that you don't have any epistemic reasons to believe that the speaker (your friend) cares about the well-being of the third friend in the way that you do. When deciding whether to believe the speaker, you need to determine what will be more harmful, accepting or refusing the testimony under consideration. If the harm of refusing the speaker's testimony outweighs the harm that would be caused to the third friend if you accepted it, then it's plausible that you should accept their testimony. Your *pro tanto* moral reasons for accepting the testimony outweigh your *pro tanto* moral reasons to refuse it. But if accepting the speaker's testimony will harm your third friend more than refusing it would harm the speaker, then it's plausible that you should refuse the testimony. Now, it's probable that the harm of refusing a friend's testimony when they offer it you will outweigh the harm caused by accepting some disparaging testimony about an absent third party. In such circumstance, you have overwhelming moral reasons for accepting their testimony. Then later, when the conversation is over and you're alone, you can consider whether you still accept the speaker's testimony, given how doing so might harm your third friend. This, it seems, allows you to navigate this particular moral bind.

Finally, it's possible that there are cases in which your moral reasons are tied for and against accepting the speaker's testimony. Frankly, it's very difficult to say what you should do here. We can hope that such circumstances don't arise often—I suspect that they don't—but this hardly solves the moral problem. Accordingly, I would tend towards a pessimistic response to such situations. It's plausible, I suggest, that there are situations in which there's no acceptable course of action for you. Sometimes, unfortunately, you can't avoid hurting someone. In cases of divided loyalties, when push comes to shove, somebody will get hurt.

§6 Summary

We should treat our friends with limited epistemic partiality, or so I've argued in this chapter. I started by considering the nature of friendship. Friendship is a caring relationship, according to a mainstream philosophical conception of friendship. This means that we have *pro tanto* moral reasons to protect and promote our friends' well-being. Then I defended several ways of understanding the possibility that beliefs can harm. Since we care about our friends, and it's possible that our beliefs can harm them, it follows that we have *pro tanto* moral reasons to protect them from that harm. Epistemic partiality at the levels of evidence gathering and evidence processing enables us to protect our friend from such harm. There are, however, limits to the epistemic partiality with which we should treat our friends.

I've discussed two sources of limits to the extent to which we should treat our friends with epistemic partiality. In the chapters that follow, I'll discuss some other limitations to the extent to which we should treat our friends with epistemic partiality. In the next chapter, I'll consider whether there are epistemological reasons not to treat our friends with epistemic partiality. Then, I'll consider whether there are moral reasons not to treat our friends with epistemic partiality. Finally, I'll consider whether friendship itself is incompatible with epistemic partiality.

Chapter Three

An epistemological objection to treating friends with epistemic partiality

§1 Introduction

Taking our friends' interests into account when we form beliefs about their character and conduct might make us good friends, but it's epistemically irrational, or so some think. In fact, according to an objection I'll consider in this chapter, the costs of such epistemic irrationality outweigh whatever moral reasons we might have for treating our friends with limited epistemic partiality.

This is the epistemological objection to epistemic partiality in friendship. In this chapter, I discuss the epistemological objection but ultimately come to reject it. First, I'll introduce the objection, which is owed in part to Arpaly and Brinkerhoff (2018). Then I'll introduce and briefly explain a possible response to the objection, namely, the response that epistemic partiality is, in fact, epistemically rational. This line of response is owed variously to Kawall (2013), Schroeder (2018a), and Goldberg (forthcoming). However, I'll set this response to the objection aside in order to pursue an alternative line of defence. In the final section of the chapter, I'll argue that we should treat our friends with limited epistemic partiality *even if* it's epistemically irrational to do so. The reason, I suggest, is that the epistemological argument fails to establish that it's all things considered irrational to treat our friends with limited epistemic partiality.

§2 The epistemological objection

If treating our friends with epistemic partiality is epistemically irrational then, according to the epistemological objection, that counts against treating our friends with epistemic partiality. In fact, so the objection goes, whatever moral reasons we may possess for treating our friends with epistemic partiality, they're outweighed by the epistemic reasons we possess against epistemic partiality in friendship.

Arpaly and Brinkerhoff (2018) present this very objection.⁴⁶ They concede that we might occasionally possess *pro tanto* moral reasons to treat our friends with epistemic partiality, but argue that this shouldn't motivate any radical philosophical theory-building about what we owe our friends, epistemologically speaking (2018, 44-45). According to their epistemological objection, whatever reasons we have for treating our friends with epistemic partiality, they're likely outweighed by our *pro tanto* epistemic reasons against giving them such treatment. Epistemic partiality in friendship is over-rated, they claim, because those defending the position have over-estimated the value of irrational belief and underestimated the costs of epistemic irrationality (2018, 47). To the extent that epistemic irrationality is undesirable or bad, we possess *pro tanto* reasons not to treat our friends with epistemic partiality, or so the thought at the core of the epistemological objection goes.

Now let's look more closely at the objection, which has the following form:

- (1) Epistemic partiality in friendship is epistemically irrational.
- (2) If epistemic partiality in friendship is epistemically irrational, then it's all things considered irrational to treat our friends with epistemic partiality.
- (3) Therefore, it's all things considered irrational to treat our friends with epistemic partiality.

In sections 3 and 4 of this chapter, I'll consider why the supporters of the epistemological objection might endorse each of the premises of this argument. Then, in sections 5 and 6, I'll offer some points of rebuttal to each of these objections. The question that arises first, then, is the following: Is epistemic partiality in friendship epistemically irrational?

§3.1 Epistemological concerns

Perhaps we should be concerned about the epistemic rationality of the beliefs and judgements that we come to hold as a result of epistemic partiality in friendship. After all, there is *something* unsettling about allowing apparently irrelevant factors to determine our beliefs. For instance, we might be worried about the epistemic rationality of our moral convictions or philosophical

⁴⁶ In fact, Arpaly and Brinkerhoff (2018) consider two versions of epistemic partiality. At first, they consider whether treating a friend with epistemic partiality is like taking a counter-evidential belief pill—that is, a pill that makes you believe regardless of your evidence. The first chapter of this thesis suggests that this is not a good picture of epistemic partiality in friendship.

beliefs, if we found out that we only held them because of how we were raised or where we studied.⁴⁷ Likewise, we may have concerns about the epistemic rationality of those beliefs which we only hold because of our personal connections with others. Describing these kinds of concerns, Paul and Morton write that

[w]e normally think of allowing our beliefs to be influenced by our personal relationships and emotions as the paradigmatic case of epistemic irrationality, insofar as our relationships and emotions are irrelevant to the truth of those beliefs. (Paul and Morton 2018, 76)

The fact someone is your friend doesn't make a difference to whether your evidence supports a favourable or a damning judgement about their moral character and conduct. Accordingly, the fact that someone's a friend shouldn't, epistemologically speaking, exert an influence over what conclusions you draw from a given body of evidence about them. When the fact that you're friends with someone does, in fact, make a difference to what conclusions you draw from the evidence you possess about them, so the thought goes, you're epistemically irrational.⁴⁸

Surprisingly, the claim that treating your friends with epistemic partiality is epistemically irrational isn't denied by the most prominent defenders of epistemic partiality in friendship. A quick survey of the literature is telling. For Stroud, it's the epistemic irrationality of epistemic partiality in friendship that distinguishes it from epistemic impartiality. She writes the following:

Friendship positively demands epistemic bias, understood as an epistemically unjustified departure from epistemic objectivity. Doxastic dispositions which violate the standards promulgated by mainstream epistemological theories are a constitutive feature of friendship. Or, to put the point as succinctly—and brutally—as possible, friendship requires epistemic irrationality. (Stroud 2006, 518)

If it's a constitutive feature of friendship that we form beliefs in violation of the norms of mainstream epistemological theories, then, assuming one of those theories is correct, we should

⁴⁷ Cohen describes these worries in his chapter, 'Paradoxes of Conviction' (2000).

⁴⁸ Of course, it's not always worrying to learn that our relationship with someone has made a difference to what we believe about them. For example, when our friends tell us their secrets, and they do so because they're our friends, our relationship has made a difference to what we believe. But we don't worry that the beliefs we come to hold after a friend confides a secret in us are epistemically irrational.

worry about the epistemic status of the beliefs we form about our friends. Alternatively, you might suspect, along with Stroud, that this suggests that those mainstream epistemological theories are false.

Stroud isn't alone in thinking that the beliefs we form about our friends are often epistemically irrational, or at least, that they are according to mainstream epistemological theories. According to Keller, when we treat our friends with epistemic partiality, what we come to believe is determined not by truth-relevant factors, but rather, by factors that have to do with their well-being. What's more, that's constitutive of what it is to be a good friend.

[W]hen good friends form beliefs about each other, they sometimes respond to considerations that have to do with the needs and interests of their friends, not with aiming at the truth [...] The fact that a person is someone's friend can sometimes explain why that person is inclined to believe certain sorts of falsehoods. (Keller 2004, 330)

For Keller, part of what it means to be a good friend is for your beliefs and other doxastic attitudes to be responsive to factors that are not truth-oriented. The implication, so it seems, is that this leads us further from the truth, and that this is epistemically irrational.

§3.2 Is epistemic partiality in friendship epistemically irrational?

Stroud (2006) and Keller (2004, 2018) employ what Hawley (2014, 2039) calls a 'counterfactual test of [epistemic] reasonableness'. The counterfactual test purports to show that epistemic partiality is epistemically unreasonable by comparing the beliefs we come to hold about our friends given a certain body of evidence and the beliefs we would form about non-friends, given the same total evidence.⁴⁹ Unfortunately, despite its simplicity, the counterfactual test doesn't show that epistemic partiality is epistemically irrational. As Hawley explains (2014, 2039), even if it's true that we come to hold different beliefs about our friends than we would about non-

⁴⁹ Evidentialism is the view that our beliefs are justified if and only if they're supported by our evidence (see Feldman and Connee 1985 for an influential defence of the view and Dougherty 2011 for an extended critical discussion of it). Note that, on the account of epistemic partiality presented in Chapter One, epistemic partiality is either 'upstream' of purely evidentialist concerns or 'downstream' of them. At no point is the suggestion that we should believe regardless of our evidence—epistemic partiality is a question of how we respond to our evidence, rather than whether we respond to it.

friends given the same evidence, that doesn't show that our beliefs about our friends are the epistemically irrational ones.

To support the claim that epistemic partiality at the level of evidence processing is epistemically irrational, it must be shown that epistemic partiality in friendship involves the application of epistemically irrational epistemic standards. In the following section, I'll offer some reasons for thinking that when we treat our friends with epistemic partiality, we apply epistemically irrational epistemic standards.

Now, in Chapter One I argued that epistemic partiality involves bias at both the level of evidence processing and evidence gathering. So far, the discussion has focussed on the epistemic (ir)rationality of evidence processing in epistemic partiality. If it's the case that epistemic partiality at the level of evidence processing is epistemically irrational, that would be sufficient for the purposes of the epistemological objection. Nevertheless, there are some reasons to think that epistemically partial evidence gathering is also epistemically irrational. I'll explain these reasons in due course.

§3.3 Epistemically irrational evidence gathering

Why might a supporter of the epistemological objection think that epistemic partiality in friendship is epistemically irrational? Epistemic partiality in friendship at the level of evidence gathering is epistemically irrational if it involves the application of epistemically irrational evidence gathering processes.

Whether an evidence gathering process is epistemically rational is a function of whether it serves our epistemic goals. Understood as such, epistemic rationality is a branch of instrumental rationality, except that it focusses exclusively on serving our epistemic goals rather than on our practical ones. And what are our epistemic goals? Well, if this view of epistemic rationality is correct, then holding true beliefs and not holding false beliefs are our primary epistemic goals.⁵⁰

⁵⁰ *'We must know the truth; and we must avoid error,— these are the first and greatest commandments of would-be knowers'*, at least according to William James (1897/1992, 469). We must be careful how we interpret this. As Piller (2009) argues, someone who's interested in believing that *p* only if *p* is true isn't concerned with how the world is, but rather, with being right about the world. Some epistemologists (e.g. Timothy Williamson 2000) think that knowledge is the norm of belief; that's consistent with my view, since knowledge is factive and, therefore, a knowledge

We can start to see this, so the thought goes, by considering the nature of epistemic evaluation. ‘Epistemic evaluation’, writes William Alston,

is undertaken from what we might call “the epistemic point of view.” That point of view is defined by the aim at maximizing truth and minimizing falsity in a large body of beliefs. . . . For a belief to be justified is for it, somehow, to be awarded high marks relative to that aim. . . . Any concept of epistemic justification is a concept of some condition that is desirable or commendable from the standpoint of the aim at maximizing truth and minimizing falsity . . . (Alston 1985, 83-84.)

Note that on Alston’s view epistemic justification is ultimately concerned with seeking truth and avoiding error. The reason we seek epistemically justified beliefs is that they’re more likely to be true (and less likely to be false) than beliefs that lack epistemic justification. Laurence Bonjour, similarly, says,

The distinguishing characteristic of epistemic justification is thus its essential or internal relation to the cognitive goal of truth. It follows that one’s cognitive endeavors are justified only if and to the extent that they are aimed at this goal, which means very roughly that one accepts all and only those beliefs that one has good reason to think are true. (Bonjour 1985, 7-8.)⁵¹

Admittedly, Alston and Bonjour are writing about epistemic justification. What does this mean for the epistemic rationality of evidence gathering processes, if the acquisition of true beliefs and the avoidance of false ones are our epistemic goals? An evidence gathering process is epistemically rational to the extent that it serves these epistemic goals, which is to say that it increases the chances of our believing the truth and decreases our chances of believing falsehoods.

With these goals in mind, we’re in a better position to consider what it would mean for epistemic partiality at the level of evidence processing to be epistemically irrational. Epistemic partiality in friendship is epistemically irrational if it is inimical to our epistemic goals. It is

norm for belief logically entails the truth norm of belief. Kvanvig (2005) argues for pluralism about the aim(s) of belief.

⁵¹ Marian David (2001, 151) refers to these passages as paradigmatic examples of the invocation of the ‘truth-goal’.

inimical to our epistemic goals if it makes us less likely to believe truthfully and more likely to believe falsely.

It's plausible that we already embrace the notion that there are epistemically good and bad ways of gathering evidence. For example, if someone gets all their information about politics from unreliable websites, we'd likely doubt that they know what they're talking about, at least on political matters. But if someone thoughtfully considers opposing views to their own and carefully checks the credibility of the sources they have relied on, then we might think that they're more likely to get to the truth of the matter. In this section, I'll briefly consider what it means for an evidence-gathering process to be epistemically irrational. Then I'll suggest why someone who supports the epistemological objection might think that epistemic partiality at the level of evidence gathering might be epistemically irrational.

Let's consider a simple example in which someone's evidence gathering practices appear to be irrational. Sylvia has a cold. She's heard that some people think the healing power of crystals can cure her cold. She's not sure, so she decides to look on the internet for more information. She quickly searches on Google for the following phrase: 'Crystal healing cure cold'. The first twenty results suggested by Google are supportive, in one way or another, of the claim that crystal healing can cure her of the common cold. So, she comes to believe that crystal healing can cure her cold based on the evidence she acquired through her Google search.

Meet Héctor. Héctor is a fanatical follower of a populist, far-right politician, Jaime. Among other things, he views Jaime as an unimpeachable moral leader. To begin with, Héctor didn't know much about politics or economics. Rather, he was seduced by the spirit of Jaime's brazen nationalism. It appealed to his passions rather than his reason. It's for these reasons that he originally came to support Jaime. However, there is overwhelming, publicly available evidence that Jaime is corrupt. Héctor, however, will not countenance the idea that Jaime is corrupt. In fact, he refuses to pay any attention to the publicly available evidence of Jaime's corruption. As soon as the topic of Jaime's corruption comes up in conversation, he changes the subject or ends the conversation altogether. He won't read any news articles from newspapers that are critical of Jaime's government. However, Héctor spends a great deal of time reading articles on far-right leaning websites where the allegations of corruption against Jaime are frequently dismissed as 'fake news'.

Sometimes people seem irrational because of the way in which they pay attention to certain features of an experience. Now let's consider a more complicated example, owed to Susanna Siegel (2017, 158). The setting is a job interview. The interviewer is prejudiced against members of a particular ethnic minority. When the interviewer evaluates the interviewees'

responses to each of their questions, they pay much more attention to any traits that they perceive as negative. The interviewer is, if not trying to trip up the interviewee, at least eager to note any sign of unsuitability for the job. When the interview is complete, the interviewer has a very negative overall view of the interviewee's suitability. The interviewer hasn't fabricated negative traits, but she's gathered evidence in a way that gives a more prominent place for evidence of negative traits in her total evidence.

What should we make of these cases? In each of these cases, the subject employs epistemically irrational evidence-gathering processes. To understand why these processes are epistemically irrational, we must consider how they, as evidence-gathering practices, fail to serve their epistemic goals. In the crystal-healing case, Cynthia sincerely wants to find out whether crystals will cure her cold. However, the method of evidence-gathering that she employs does not serve this purpose well at all. Perhaps all that her evidence really rules out is the possibility that there's unanimous agreement that crystal healing will not cure her cold. In the case of Héctor's political convictions, Héctor doesn't sincerely want to find out whether Jaime is in fact corrupt. Rather, he wants to protect his political hero. By seeking only evidence that supports his pre-existing beliefs about Jaime and refusing to gather or engage with evidence that could challenge this view, he reduces his chances of finding out the truth about Jaime's corruption, and increases his chances of continuing to believe falsely that the allegations of corruptions are untrue. Finally, in the interview case, adapted from Siegel (2017), the interviewer may sincerely wish to find out whether the candidate is suitable for the job. But the way they focus on the negative attributes of the ethnic minority candidate does not serve this goal. The interviewer is less able to determine whether the interviewee is a suitable candidate as a result of their biased evidence gathering: By focussing on the ethnic minority candidate's negative traits, the interviewer gathers evidence in a way that makes it much less likely that they will form an accurate judgement about the interviewee's suitability.

In each of these cases, we find evidence-gathering processes that seem to be epistemically problematic. According to the epistemological objection, these processes are epistemically irrational, because they fail to serve our epistemic goals of believing the truth and avoiding error. Now, the next step for the objector is to show that similar kinds of epistemically irrational evidence-gathering processes are involved when we treat our friends with epistemic partiality. So, does epistemic partiality in question involve epistemically irrational evidence-gathering processes?

The objector argues that it does. Consider the example of *The Third Man*. Holly Martins gathers evidence about Lime in a very particular way. From the offset, Martins rejects the

allegations against Lime. He sets out not to find out the truth about Lime, but rather, to seek proof of Lime's innocence. When he encounters evidence which supports his belief that the allegations against Lime are true, he accepts it more-or-less unquestioningly. When he encounters evidence which doesn't support his conviction about Lime's innocence, but rather, challenges it, he expends a great deal of cognitive effort trying to explain away or else diminish the importance of that evidence.

In Keller's Poetry Reading example (2004), Eric gathers his evidence about Rebecca in a very particular way too. Eric pays much more attention to the positive features of Rebecca's poetry reading than he does to its negative qualities. What evidence he ends up with is, to echo Siegel (2017, 159), a function of what features of her poetry reading he pays attention to. By paying more attention to positive features of Rebecca's poetry reading, Eric will tend to gather more evidence that supports a favourable view of her qualities as a poet. The outcome of this process is that a greater proportion of Eric's evidence will support the favourable view of Rebecca's poetry skills. This is epistemically irrational, because it does not bring him closer to finding out the truth about Rebecca's poetry reading, and it increases the chance of him being mistaken.

For these reasons, according to the objection, we should think that epistemic partiality at the level of evidence gathering is epistemically irrational. Now, it would be sufficient for the objector's purposes to show that only one component of epistemic partiality in friendship was epistemically irrational. However, the objector can strengthen their case by arguing that epistemic partiality is epistemically irrational at the level of evidence processing as well as at the level of evidence gathering. I'll attempt to present that argument in the following section.

§3.4 Epistemically irrational evidence processing

People who apply overly lax or overly strict epistemic standards when they process their evidence sometimes come in for epistemic criticism. This criticism is epistemological in character. When we complain that someone is being too gullible or too sceptical, we're complaining that they too easily come to believe falsehoods or that they're too resistant to the truth that lies within their reach. Given the conception of epistemic rationality considered above, we can see how it could be epistemically irrational to apply overly strict or overly lax epistemic standards. Overly lax ones increase our chances of forming false beliefs. Overly strict epistemic standards prevent us from believing the truth. Thus, according to the epistemological objection, when people fail to apply *rational* epistemic standards, they process their evidence in an

epistemically irrational way. (This claim will come under scrutiny in section 5.) Before we consider what this means for epistemic partiality at the level of evidence processing, let's consider a couple of neutral examples. Then I'll suggest an explanation of why supporters of the epistemological objection could conclude that epistemic partiality at the level of evidence processing is epistemically irrational.

Imagine a textbook case of gullibility. Aaron is usually a very discerning judge of testimony, not when it comes to his co-worker, Bradley. Aaron believes everything Bradley tells him. In fact, Bradley is not a reliable source of testimony. Bradley is a serial liar and a bullshitter too. He lies about the little details of his days and he pretends to know about all sorts of things when he really has no idea what he's talking about. What's more, this is all quite obvious—at least, it's obvious to most people. If Aaron applied his usual epistemic standards when he considered Bradley's testimony, he'd never believe anything Bradley said. But for some reason, Bradley has a strange effect on Aaron, such that he doesn't evaluate Bradley's testimony in the discerning way that he'd evaluate anyone else's. Rather, he applies much lower standards. As a result, Aaron comes to form all sorts of false beliefs.

Now let's consider another example. Cristina thinks that she should believe what scientists tell us about climate change. She came to that view after learning that scientists were probably in the best position to know about climate change and that, even though they're fallible, they still represent her best option for finding out about such matters. However, she won't accept what scientific consensus tells her about vaccines. When it comes to the question of the safety and efficacy of vaccines, she sets considerably higher standards. In fact, her standards are so high that it's unlikely she'll ever believe what scientific consensus says about vaccines. She's a vaccine-sceptic in the sense that she thinks that we can't tell whether the findings of modern medical science are true or false. As a result, there's a valuable array of truths that she never comes to believe.

Why think these practices are epistemically irrational? One thing to note is that it's not the partiality itself that leads to epistemic irrationality. Rather, it's the irrational application of epistemic standards.⁵² What's irrational about the application of higher or lower epistemic standards in these cases is that the subject is led further from truth and closer to error. Granted the view of epistemic rationality outlined above, we should apply our epistemic standards in ways

⁵² Whether the application of this-or-that epistemic standard is epistemically rational depends on whether permissivism is true. The implications of permissivism will be discussed later in this chapter.

that best serve our epistemic goals. But by applying extremely low epistemic standards when evaluating Bradley's testimony, Aaron comes to form lots of false beliefs—many more than he would if he applied his usual standards. Similarly, when Cristina's application of very high epistemic standards when it comes to the science of vaccines has the result that she fails to form true beliefs even though the evidence in her possession would support her holding those beliefs. In each of these cases, then, the application of differential epistemic standards leads the subject away from the truth and, in the case of Aaron, towards error. Thus, the way that their evidence-processing is responsive to non-epistemic concerns (Bradley's hold over Aaron and Cristina's fear of vaccinations) thwarts their epistemic goals. Accordingly, the way they process their evidence is epistemically irrational.

Is this how we process evidence about our friends when we treat them with limited epistemic partiality at the level of evidence processing? Again, let's consider a couple of cases which lend support to the claim that it is. In the *Third Man* example and in Stroud's Sam example (2006), we find instances in which subjects withhold judgement about their friends' conduct in spite of the evidence they possess. Let's consider these examples in turn.

In *The Third Man*, it's plausible that Martins possesses sufficient evidence to believe that Lime is a criminal. For one thing, he has the testimonial evidence provided to him Major Calloway and his associates. We also know that Martins and Lime have had no contact for several years, so it's not clear that his background evidence about Lime's character would easily outweigh his new, up-to-date evidence. But rather than amend his belief, Lime suspends his judgement. In other words, he sets a much higher epistemic standard for belief with respect to the proposition that Lime is a criminal than he would for similar propositions about non-friends.

In Stroud's Sam example (2006), Sam's friends maintain their belief in Sam's good character despite the evidence they possess to the contrary. We can imagine that there are circumstances in which they don't have good epistemic reasons to reject the third-party testimony about Sam's actions, for instance, if they knew that the speaker was not usually inclined to spread rumours and had no particular vendetta against Sam.⁵³ Perhaps the third-party testimony doesn't provide enough evidence to convince them rationally that Sam really did cruelly break someone's heart, but rather, it's strong enough to warrant that they suspend their

⁵³ We might have Gricean reasons for rejecting someone's testimony (see H. P. Grice 1989). Sometimes, plausibly, it's inappropriate for a third party to tell us something insulting or critical about a friend. The inappropriateness of such utterances might provide us with reasons to reject their testimony.

judgement about Sam's good character. Their total evidence regarding Sam's conduct is inconclusive. But rather than suspend judgement, they steadfastly maintain their belief that Sam didn't, and indeed, wouldn't, knowingly break someone's heart after a one-night stand.

In each of these cases, the way the subjects process their evidence thwarts their epistemic goals of believing the truth. The truth about Lime is easily available to Martins, but he eschews it in order to protect him. And perhaps the truth isn't available to Sam's friends, not yet, but even so, their loyalty to Sam leads them to risk believing falsely that the third party's testimony is false when, given their inconclusive evidence, they should withhold judgement about Sam's conduct.

In summary, supporters of the epistemological objection have some more or less persuasive reasons for thinking that epistemic partiality in friendship is epistemically irrational. Of course, for the purposes of the epistemological objection, they only need show that one of the two varieties of epistemic partiality is epistemically irrational. Now, the next step for the epistemological objection is to show that it's all things considered irrational to treat our friends with epistemic partiality. That will be the subject of the next section.

§4 Is epistemic partiality irrational, all things considered?

The second premise of the epistemological objection maintains that if epistemic partiality in friendship is epistemically irrational, then it's all things considered irrational to treat our friends with epistemic partiality. There are, I think, two reasons why supporters of the epistemological objection might endorse this premise.

First is the argument from the exclusivity of epistemic reasons, or the exclusivity argument. According to the exclusivity argument, only epistemic reasons are normative reasons to believe. Exclusivity is the name of the view that only epistemic considerations are normative reasons to form certain doxastic attitudes.⁵⁴ There are not moral or prudential reasons to believe, or so the thought goes. This view is closely related to the notion of transparency, introduced by Nishi Shah (2003). Roughly, transparency says that the only relevant considerations which bear on whether you should believe that p are considerations which are relevant for determining whether p (Shah 2003, 447). This seems to entail exclusivity.⁵⁵ To reiterate: the only

⁵⁴ We must distinguish between psychological and normative exclusivity. See Conor McHugh 2015 for a discussion of the relationship between these two views. In this discussion, I'm concerned with normative exclusivity.

⁵⁵ This claim is disputed, especially by McHugh (2015).

considerations which bear on whether you should believe (or disbelieve, or withhold judgement) that p , are considerations which bear on whether p . Consequently, if it's epistemically irrational to believe that p , then it's irrational to believe that p all things considered.

Now for the key move in the exclusivity argument. If it's epistemically irrational to believe that p all things considered, then it's irrational to bring it about that you believe that p . Non-epistemic, prudential versions of this inference are intuitively appealing, at least. For instance, if it's prudentially irrational to eat only egg shells and drink only seawater, then it's prudentially irrational to bring it about that you eat only egg shells and drink only seawater. Let's think of epistemic partiality as a way of influencing your beliefs—what William Alston (1988, 275) would describe as a way of exercising 'long-range control' over them. Even if it so happens that we're able to influence our beliefs and judgements in the ways that are required for epistemic partiality in friendship, that makes no difference to whether we should. When considering whether it would harm a friend if you believed some relevant proposition that p , you don't gain reasons to believe or disbelieve that p . That's because, plausibly, in a great many cases, for all you know, the fact you have *pro tanto* moral reasons to (bring it about that you) believe that not- p doesn't entail that not- p . It's not even evidence that not- p . And if you don't gain reasons to believe that not- p , you don't gain reasons to influence your beliefs and bring it about that you believe that not- p . Your *pro tanto* moral reasons simply don't make a difference to what you should do. This is one reason why the objector might think that if epistemic partiality in friendship is epistemically irrational, then it's irrational all things considered.

Second is the no-safe-dose argument. Arpaly and Brinkerhoff claim that '[A] here is no safe dose of epistemic irrationality—once you have allowed it into your life, you are never safe from its repercussions, and neither are your friends' (2018, 47). According to the no-safe-dose argument, the costs of epistemic irrationality swamp out any *pro tanto* moral reasons we might have for engaging in epistemically irrational belief-forming practices such as those involved in epistemic partiality in friendship. This claim is reminiscent of William Clifford's claim that even small, apparently harmless transgressions against his ethics of belief are ultimately immoral. He writes the following:

Belief, that sacred faculty which prompts the decisions of our will, and knits into harmonious working all the compacted energies of our being, is ours not for ourselves but for humanity. [...] It is desecrated when given to unproved and unquestioned statements, for the solace and private pleasure of the believer; to add a tinsel splendour to the plain straight road of our life and display a bright mirage beyond it; or even to drown the common sorrows of our kind by a self-deception which allows them not only to cast down, but also to degrade us. Whoso would deserve well of

his fellows in this matter will guard the purity of his beliefs with a very fanaticism of jealous care, lest at any time it should rest on an unworthy object, and catch a stain which can never be wiped away. (1877, 292-293)

The language of desecration is powerful here. The thought is that once truth is dethroned from its rightful place as the goal of our epistemic activity, all manner of epistemic, moral, and prudential failures will surely follow. Let's consider two possible costs of desecrating belief and dethroning the truth.

Possibly, a major cost of embracing epistemic irrationality is that it's probable that the epistemically irrational beliefs we come to hold about our friends have a knock-on effect on the epistemic rationality of other beliefs we hold. This could occur when we attempt to draw further inferences from our epistemically irrational beliefs about our friends to other matters in which gaining the truth is our priority. It's all very well throwing caution to the wind and giving our friends the benefit of the doubt today, but we can't be sure that tomorrow we won't need to rely on those beliefs in situations where grasping the truth is of paramount importance.

This can have purely epistemic consequences, such as when epistemically irrational beliefs impede our access to further knowledge, for instance about our friends. But there can be moral and prudential consequences true. If, like Sam's friends, you refuse to believe that a friend is capable of acting cruelly towards a third party, you risk causing harm against that third party. If the accusation is morally grave in nature, it's plausible that you owe it to the third party to accept the accusation. (This is closely related to an objection discussed in Chapter Four.) If this is correct, then there's no safe dose of epistemic irrationality. Therefore, it's epistemically irrational to treat our friends with *any* epistemic partiality.

In summary, we've considered two reasons why the objector might endorse the claim that if epistemic partiality in friendship is epistemically irrational, then it's all things considered irrational to treat our friends with epistemic partiality.

§5.1 Response: Epistemic partiality is not epistemically irrational

Now, one way to handle the epistemological objection to epistemic partiality in friendship would be to deny that epistemic partiality is epistemically irrational. In this section, I'll introduce a couple of possible responses to the positions presented in section 3. Once I've discussed these responses, though, I'll set them aside for the sake of argument. Then, assuming that limited

epistemic partiality in friendship is epistemically irrational, I'll propose one last line of response to the epistemological objection.

§5.2 Epistemic partiality in evidence gathering is not epistemically irrational

Jason Kawall (2013) has argued that epistemic partiality at the level of evidence gathering is not epistemically irrational.⁵⁶ Rather, he suggests that the way we gather evidence about our friends is rather like the harmless way we pay more attention to things we like or that are important to us (2013, 352). He offers the following story for comparison:

We walk into a pub—as a whisky-drinker you immediately look to those bottles and form justified true beliefs about them; as a beer-drinker I fail to form the beliefs that you do, but given my interests I form justified true beliefs about what beers are available. (2013, 352.)

The thought is that it's unreasonable to conclude that the way we gather evidence about our friends is likely to lead to epistemic irrationality from the fact that being friends with someone makes a difference to the body of evidence you gather about them in otherwise equal circumstances. If Eric pays attention to different aspects of Rebecca's poem, he'll likely come to hold different beliefs about it from those an impartial observer would hold. That doesn't imply that Eric's beliefs are likely to be unjustified; indeed, the beliefs he forms may be well justified but based on a different body of evidence from the one possessed by the impartial third party.

Kawall's characterisation isn't exactly incorrect: to be sure, there will be occasions in which a friend and a non-friend come to different conclusions about a third person based on

⁵⁶ Sanford Goldberg objects in particular to the claim that epistemic partiality is epistemically irresponsible. He writes that

friendship itself is another instance of this same general phenomenon whereby we can have practical reasons to re-open inquiry: given their friendship, if S tells A something important about herself that goes against A's current evidence, A "owes it" to S to get more evidence, and in particular to make sure A has explored all of the places that are most likely to deliver further significant evidence (including evidence vindicating S's say-so). This is no more impermissible from an epistemic point of view than [...] a subject's looking for more evidence in cases involving particularly high stakes. These cases merely make manifest that we sometimes have practical reasons to re-open inquiry. (Forthcoming, 12)

different considerations, even though they (broadly speaking) witnessed the same event(s). But it doesn't generalise as broadly as he intends. It doesn't show that epistemic partiality isn't epistemically irresponsible. First, let's consider the poetry case again. Of course, it's possible that two people make different observations about her poetry, and that they therefore come to different conclusions about its quality.

Let's think about a variation of Kawall's example. Suppose you and I had a specific question in mind, namely, what do they have to drink at this bar? We can imagine situations in which this question might be pertinent.⁵⁷ Perhaps we're scouting ahead of a departmental drinks reception. Now, because I like whisky and you like beer, we gather different evidence. We come to different conclusions about what there is to drink at the bar because we gathered different evidence. Moreover, the ways each of us gathered evidence was shaped by our preferences, just like in Kawall's example. Just as I looked for whisky because I prefer bars with whisky, so you looked for beer because you prefer bars with beer. In some sense, we've both failed. Remember, the question we set out to answer was 'what do they have to drink?' but, owing to our different preferences, we gathered evidence differently. The thing is, we shouldn't have done that, epistemologically speaking. Our respective preferences for different types of alcoholic drinks are irrelevant to the question what drinks they sell in this bar. Our preferences, therefore, shouldn't make a difference to what drinks we come to believe are available at the bar. Let's return to Keller's Poetry Reading example (2004). Let's imagine that both Eric and a third party who isn't Rebecca's friend are watching her performance, but moreover, suppose that they're both considering whether it's a good poem. Suppose it's true that they come to different conclusions simply because they made different observations, and that, had they made the same observations, they'd have the same conclusions. Granting that, we can still criticise one or both of them for the way they gathered that evidence! This is perhaps even more obvious if we consider a different case, like a job trial. Suppose that Rebecca co-owns and runs a café with her business partner Elisa. One day, Eric applies for a job at the café. He passes the interview and is invited, with one other candidate, Amir, to complete a trial shift. Rebecca and Elisa observe Eric and Amir as they go about their duties on the trial shifts. At the end of the day, once they've said their farewells, Rebecca and Elisa discuss who they thought was the better candidate. Predictably, Rebecca thinks that Eric put on a better performance, but Elisa disagrees. From Elisa's perspective, Amir was better. Now, suppose that the reason for their difference of

⁵⁷ Questions are important indicators of our epistemic goals. See Christopher Hookway 2008 for a discussion of the place of questions in epistemology.

opinion is that Rebecca expended much more effort looking for the positive side of Eric's performance. She made sure to look out for signs of his suitability for the post. After all, she wants him to get it. What's more, she might have avoided making observations where she thinks there's a good chance he (or anyone) could make a mistake. She was, in a word, biased. I'll return to the subject of bias in evidence-gathering and epistemic norms after considering another example.

Secondly, when a specific concern is raised about some aspect of a person's character or conduct, then the analogy with looking at drinks in a bar breaks down. If a friend is accused of knowingly breaking someone's heart or of stealing a tip-jar, as in Stroud's example and my own, then, on the assumption that both the interested party and the disinterested party want to find out the answer to the same questions: Did Sam knowingly break someone's heart? Did John steal from the tip-jar? In cases like this, it seems that the good friend and the disinterested party both want to find out the truth, but they go different ways about it.

Kawall's account of epistemic partiality in evidence gathering doesn't capture what's wrong with Rebecca's actions here. In both of these cases, the problem isn't that they came to different conclusions because they had different evidence; rather, the problem is that they had different evidence in the first place!

§5.3 Epistemic partiality in evidence processing is not epistemically irrational

In this section, I'll consider two possible responses to the claim that epistemic partiality in evidence processing is epistemically irrational. First I'll consider Hawley's permissivist response (2014), according to which it's possible that there's more than one epistemically justified doxastic response to a body of evidence. Second, I'll consider Schroeder's pragmatic-intellectualist response (2018a), according to which it's sometimes reasonable all things considered, to withhold judgement rather than form harmful beliefs about your friends.

Perhaps it's epistemically rational to treat apply different epistemic standards when we form beliefs about our friends from those we would apply when thinking about non-friends, other things being equal. Hawley argues that if permissivism is true, then it's possible that treating our friends with epistemic partiality doesn't leave us with epistemically unjustified beliefs. The epistemological objection to epistemic partiality at the level of evidence processing rests on the assumption that there's at most one correct response to a given body of evidence, or, in other words, a body of evidence supports at most one conclusion. Feldman formalises this notion in what he calls 'the uniqueness thesis':

[A] body of evidence justifies at most one proposition out of a competing set of propositions (e.g., one theory out of a bunch of exclusive alternatives) and that it justifies at most one attitude toward any particular proposition. (2007, 205)

What does the uniqueness thesis tell us about the beliefs that we hold as a result of epistemic partiality in friendship? Let's consider the Third Man example. Martins amasses a large body of evidence regarding Lime's disappearance. The majority of the evidence Martins possesses is testimonial. It mostly relates details about Lime's movements around Vienna. Let's suppose, contrary to our earlier discussions of the case, that he shares his total evidence with Major Calloway. Despite sharing their evidence, Martins and Calloway come to different conclusions about Martins' character. While Calloway comes to a negative judgement about Lime's character and conduct based on her total evidence, Martins comes to a very different conclusion. He concludes that there's a conspiracy to frame Lime for crimes he didn't commit.

We can think of Martins and Calloway as having two, competing interpretations of the same body of evidence. What the uniqueness thesis implies, when applied to this case, is that at least one of their interpretations must be mistaken. The thought is that Martin comes to a different conclusion from Calloway because of the irrelevant influence of his friendship with Lime. But where irrelevant influences exert an influence on our belief, we should worry about that belief's rationality.

The thing is, the intuitive support for Uniqueness Thesis is limited. Hawley notes that there are just as many cases when evidence seems more permissive than the Uniqueness Thesis can allow. She writes:

Different people—and the same person in different moods—can reasonably differ in their doxastic policies. Some people are somewhat quicker to belief than others are, some put more weight on the evidence of their own senses than others do, some are more sceptical about testimony than others are. There are many reasonable attitudes to epistemic risk, many acceptable ways to strike the balance between pursuing true belief and avoiding false belief. We are all familiar with situations in which there is more than one conclusion which could reasonably be drawn, even against the same background assumptions. For example, disagreements between colleagues about what grade a student paper deserves are not always reducible to disagreements about what qualities a student paper ought to exhibit. One can be an atheist, and regard oneself as epistemically reasonable in this respect, without regarding theists as epistemically

unreasonable. Scientists can reasonably disagree about the merits of a given theory, even though all have access to the relevant evidence. (2014, 2040)⁵⁸

Perhaps we should reject the uniqueness thesis and instead adopt the position that evidence is permissive. Among the supporters of permissivism, there's an ongoing disagreement about how just how permissive a body of evidence can be.⁵⁹ I'll work with the assumption that extreme permissivism is true: It's possible that a body of evidence support belief that p and that not- p at the same time.

Miriam Schoenfield (2014) suggests that permissivism might help to resolve the problem of irrelevant influences on belief. The problem is that irrelevant influences appear to explain why different people come to different conclusions based on the same body of evidence. It appears that irrelevant influences play a significant role in determining a great deal of our moral and religious convictions, to say the least. If, as the thought goes, irrelevant influences threaten the rationality of a belief, their prevalence in the causal history of our actual beliefs gives rise to sceptical concerns.

Here's where permissivism comes in. If Martins' evidence is permissive, then it's possible that it's epistemically rational to apply stricter epistemic standards when he forms judgements about Lime's conduct than, for instance, Calloway would apply to the same evidence. Indeed, if permissivism is true *and* the evidence is permissive, then Limes and Calloway could come to form different beliefs about whether Lime is a criminal, and both would be epistemically rational

⁵⁸ Gideon Rosen, in a passage quoted by Miriam Schoenfield (2014, 196), speaks in favour of the intuitive appeal of permissivism about evidence. He writes:

It should be obvious that reasonable people can disagree, even when confronted with a single body of evidence [...] Paleontologists disagree about what killed the dinosaurs. And while it is possible that most of the parties to this dispute are irrational, this need not be the case. To the contrary, it would appear to be a fact of epistemic life that a careful review of the evidence does not guarantee consensus even among thoughtful and otherwise rational investigators. (2001, 71)

⁵⁹ Extreme permissivism allows that there are bodies of evidence which support belief in p and in not- p . Moderate permissivism, on the other hand, allows only a narrower range of responses to a given body of evidence could ever be permissible. I'll not attempt to resolve this disagreement here. Sophie Horowitz (2014) argues persuasively that of extreme permissivism, moderate permissivism, and the uniqueness thesis, moderate permissivism faces the most serious challenges.

in maintaining those beliefs, even though they share their evidence. Likewise, in Stroud's Sam example (2006), if Sam's friends' evidence is permissive, then it's possible that even though they come to a different conclusion about Sam's conduct than a disinterested party would on the same evidence, both they and the disinterested party would be epistemically rational in maintaining their respective judgements.

If permissivism is true, then at least some cases of epistemic partiality at the level of evidence processing do not seem to be epistemically irrational after all. Now I'll consider another response to the claim that epistemic partiality is irrational, owed to Mark Schroeder. Schroeder (2018a) argues that epistemic partiality in friendship is not epistemically irrational. He explains this in terms of pragmatic intellectualism. Pragmatic intellectualism is the view that it's practically rational to withhold from judgement when the risks of holding a false belief outweigh the risk of not forming a judgement.

Schroeder's pragmatic intellectualism is a development of existing pragmatic encroachment views. Pragmatic encroachment views share in their support for one key claim: Pragmatic considerations make a difference to whether your evidence provides epistemic justification for your belief, or whether your belief counts as knowledge.⁶⁰ Pragmatic encroachment views are often motivated by the way they chime with our intuitions about the relation between knowledge attribution, or justified-belief attribution, and what's at stake, practically speaking. Let's start by considering a pair of such cases:

Low stakes: Did you lock the door to your apartment behind you? Well, you usually do and you faintly recollect locking the door. Although sometimes you can't rely on your memory, especially, when it comes to habitual actions like locking the door, you continue to hold the belief that you did lock up on your way out. After all, you're only popping to the shop to pick up snacks for your big trip and your partner is still at home, so it doesn't really matter if you're mistaken.

High stakes: You've just left your apartment for the second time. This time, you're going on holiday for two weeks. You're just waiting for the taxi to the airport. 'Did I lock the

⁶⁰ Jeremy Fantl and Matthew McGrath (2002, 2009), John Hawthorne (2003), Jason Stanley (2005), Stanley and Hawthorne (2008), Schroeder (2012), Jacob Ross and Schroeder (2014), among others, have defended a variety of pragmatic-encroachment-type views. See Blake Roeber 2018 for an insightful critique of the debate as it stands.

door?’ You ask yourself. At first you’re fairly confident: you usually lock it. But then you’re struck by a pang of doubt. It occurs to you that you can’t always rely on your memories when it comes to habitual actions like locking the door. What’s more, your neighbour’s apartment was burgled last week. You can’t afford to make the same mistake! You suspend judgement and rush back to check the door is firmly locked.

In Low Stakes, it seems like your belief that you locked the door is epistemically justified. However, in High Stakes, it seems like your belief isn’t justified. Though you have the same evidence (i.e., you remember locking the door and you think you have a reliable door-locking habit), there’s one crucial difference between the cases. In the first case, the stakes are much lower. You have nothing to lose if it turns out that you were mistaken. In the second case, the stakes are much higher. You stand to lose an awful lot if your apartment is burgled and you were mistaken about whether you locked the door. One very simple conclusion we might draw from this is the following: The higher the possible costs of believing that p when p is false, the more evidential support you need for your belief that p to be justified. Evidential support is necessary for justified belief, but pragmatic considerations determine how much evidential support is needed.

In the examples provided above, what you know makes a difference to the rationality of what you do. You have the option to go back and check the door is locked or continue and hope for the best. In the low-stakes case, it seems that it’s rationally permitted for you to continue to the shops without checking whether you locked the door. In the high-stakes case, it seems that it’s not rationally permitted for you to continue to the airport. The notion that what you know makes a difference to the rationality of what you do has been formalised in so-called Knowledge Action principles.

If evidentialism is correct then epistemic partiality in friendship may well lead us to form unjustified beliefs, or to withhold from forming justified beliefs which are in some sense available to us. However, as Kawall notes, ‘a growing number of epistemologists would reject a strict division between epistemic and pragmatic norms’ (2013, 366). If some pragmatic encroachment view is correct, then it’s epistemically reasonable that a friend will require more evidence before they can believe with justification something that could harm their friend.

Kawall writes:

It is epistemically required that the good friend seek greater evidence, treat a wider range of possibilities as relevant, and so on in order to know. As such, there would be no conflict between

epistemic norms and any belief-forming behaviour that is required (or merely permissible) by the standards of friendship. (2013, 366-7)

Schroeder defends a very similar view, which he calls ‘pragmatic intellectualism’. He writes:

According to pragmatic intellectualism, a belief is epistemically rational just in case the epistemic reasons to have it are at least as good as the epistemic reasons against it, evidence is the only kind of epistemic reason in favor of a belief, but in addition to evidence against its content, there are further, non-evidential, epistemic reasons against belief, among them *stakes-related* reasons, which turn on the costs of *false belief*. (2018a, 119)

Pragmatic considerations at least give us pragmatic reasons to withhold from belief, or so the thought goes. Withholding from belief is different from believing or dis-believing. As Schroeder notes in an earlier work, when you withhold your belief you abstain from making up your mind about whether *p*. He writes:

[A] natural place to look for reasons to withhold is in the costs of error. When you form a belief, you take a risk of getting things wrong that you don't take by withholding. In contrast, when you withhold, you guarantee that you miss out on getting things right. So plausibly, one important source of reasons to withhold will come from the preponderance of the cost of having a false belief over the cost of missing out on having a true belief. (2012, 277)

Now, as Schroeder writes in the passage above, typically, when we think about the kinds of non-evidential, pragmatic considerations that could count against belief that *p*, we think about the consequences of our actions that we'd undertake, if we believed that *p*. Consider the cost of error in the high-stakes example from the beginning of this section. If you leave your apartment without checking whether you locked the door because you believe that you locked it, but you're mistaken—the door is unlocked—then you risk losing a lot. If you're lucky the burglars will leave that dusty old photo album from your wedding, but they'll probably take all the expensive electronic items that you've left unguarded at home, not to mention the material damage they might cause to your furniture. As we can see, the costs of error in this case are high, so you have a pragmatic reason to withhold the belief that your door is locked until you've checked and double-checked it. In the low-stakes example, the costs of error are negligible, and as a result, you don't have pragmatic reasons for withholding belief. To do so might be considered over-cautious.

The central insight of this approach is that it's not just the consequences of our actions that provide stakes-related considerations for withholding belief. Rather, as we see in the examples discussed above, the consequences of our beliefs can too. Schroeder continues:

In the normal course of things, the costs of error which count against a belief are the consequences of the things that you will *do*, if you rely on this belief. Since one of the things that you will do if you believe a sandwich is made from almond butter is to treat it as a live option for lunch, it is clear why someone who is allergic to peanuts will face bad consequences if they falsely believe a sandwich is made of almond butter—since if that is false, it is most likely made of peanut butter, which will give them a terribly bad turn. (2018, 120)

However, as was discussed in Chapter Three, what you believe about someone has consequences for them even if your belief doesn't make any difference to your actions. As I argued, it's plausible that your beliefs about someone can harm them. In Stroud's Sam example (2006), Sam's friends would harm him if they make a negative judgement of his character. This would be very bad for them too—Sam's well-being, I've argued, is an important concern of theirs. Therefore, for Sam's friends, in this case, the cost of error are high. If pragmatic intellectualism is correct, then Sam's friends have pragmatic reasons to withhold belief when the third party tells them about Sam's misconduct.

We can find similar features in the Third Man example. Martins possesses evidence to the effect that his friend, Lime, is a murderous criminal. But if Martins comes to believe falsely that Lime is a racketeer, he stands to lose a lot. It's plausible that his belief would harm Lime. The stakes, therefore, are considerably higher for Martin than for, say, Lime's old Viennese landlady. Plausibly, the landlady doesn't stand to lose a cherished friendship if it turns out she erroneously believes that Lime is a racketeer. Indeed, it might be prudential for the landlady to be very cautious about who she allows to live in her house. What do the differences in stakes for Martins and the landlady mean for the rationality of their doxastic attitudes? Martins has prudential reasons to withhold from forming the belief that Lime is a racketeer which Lime's former landlady doesn't possess. Therefore, according to pragmatic intellectualism, Martins is epistemically justified in withholding judgement about Lime's criminal activity, but Lime's landlady is not.

If pragmatic intellectualism is correct and Schroeder's right that the limited epistemic partiality that's sometimes owed to our friends is an instance of pragmatic intellectualism in

action, then it's plausible that epistemic partiality at the level of evidence processing is epistemically rational after all.⁶¹

§5.4 Possible responses in defence of the epistemological objection

How might supporters of the epistemological objection respond to the points of rebuttal proposed in the previous sections? I'll consider the permissivist response first and then the response from pragmatic intellectualism.

The problem is that permissivism about evidence faces strong objections.⁶² Sinan Dogramaci and Sophie Horowitz defend inter-personal uniqueness on the grounds that only inter-personal uniqueness explains the practical value of the 'social practice of epistemically evaluating one another's beliefs' (2016, 131). Dogramaci and Horowitz (2016) propose that there are two theoretical virtues of inter-personal uniqueness. The first theoretical virtue of inter-personal uniqueness is that it allows us to serve as each other's epistemic surrogates (2016, 136-138): If everyone is bound by the same epistemic standards, then, if a known-competent speaker tells you that p , then you have good epistemic reasons to believe that you would have reached the conclusion they have reached, had you been in their position. If inter-personal uniqueness is false, you cannot be sure that you would have come to the same conclusion as the testifier, based on their evidence. This, surely, is an impediment to the social practice of sharing knowledge through testimony. The second theoretical virtue of inter-personal uniqueness is that it allows for the division of epistemic labour (2016, 138-139): First, if inter-personal uniqueness is true, then we're better able to divide the epistemic labour of gathering evidence. If we know that our collaborators are bound by the same epistemic standards as us, then we can send them to gather evidence and, when they return with testimonial evidence about the world, be confident that we would believe what they believe, if we'd gathered the evidence for ourselves (Dogramaci and Horowitz 2016, 138). What's more, if inter-personal uniqueness is true, then we're better able to divide the epistemic labour of reasoning. Groups of who operate according to the same epistemic standards, they write, 'can also more efficiently accomplish the task of drawing

⁶¹ Relatedly, Paul Faulkner (2018) argues that since it's possible for you to trust your friend in spite of the evidence, it's possible to believe them. And if it's possible to believe them, he argues, it's also possible to gain knowledge from their testimony.

⁶² See Roger White 2005; 2013 for other, influential responses to permissivism.

inferences on the basis of known premises' (2016, 139). That's because we can rely on our collaborators to implement the same rules of belief-formation as us.

Dogramaci and Horowitz (2016) think that inter-personal uniqueness is a plausible thesis about evidential support. Where does this leave us? Well, if interpersonal uniqueness is true, then your evidence for a proposition either does or doesn't support the doxastic attitude that you hold to it, regardless of what epistemic standards you happen to apply. The point, crucially, is that shifting your epistemic standards won't make a difference to whether your belief is justified on the basis of your evidence. Rather, it will only make a difference to what you belief hold. It seems that the main problem with the permissivist response to the claim that epistemic partiality at the level of evidence processing is epistemically irrational is permissivism itself.

The objector may attempt to rebut the response from pragmatic intellectualism too. Now, if pragmatic intellectualism is correct *and* they successfully explain the epistemic rationality of epistemic partiality in friendship, then all the better the position defended in this thesis. But there are some reasons to doubt that pragmatic intellectualism can successfully explain the epistemic rationality of all cases of limited epistemic partiality in friendship.

Pragmatic intellectualism can't explain the epistemic rationality of all instances of epistemic partiality at the level of evidence processing because it only explains the epistemic rationality of withholding judgement in spite of your evidence. It cannot explain the epistemic rationality of believing on insufficient evidence or disbelieving in spite of the evidence. But this is exactly what we see in some of the cases discussed in this thesis.

As I argued in Chapter Two, there are cases when suspending judgement isn't enough to protect or promote the well-being of your friends. Rather, in some cases, friendship specifically calls on you to believe certain propositions or disbelieve others. It's plausible, I think, that in Keller's Poetry Reading example, for instance, Eric owes more to Rebecca than just to suspend judgement about the quality of her poetry. To recall Basu's (2019) account of how beliefs can harm, Eric's failure to see Rebecca as she sees herself is what harms her, and suspending judgement about important aspects of her practical identity is hardly likely to mitigate this harm.

In summary, there are a few points of rebuttal available to supporters of the epistemological objection to epistemic partiality in friendship. I haven't discussed these in great detail. In the next section, our focus turns to the second premise of the epistemological objection.

§6 Response: Epistemic partiality is not irrational, all things considered

Let's concede, for the sake of argument, that epistemic partiality is indeed epistemically irrational. Our defence of limited epistemic partiality in friendship needn't end there. Let's consider what might be said in response to the second premise of the epistemological objection, which says that if epistemic partiality in friendship is epistemically irrational, then it's all things considered irrational to treat our friends with epistemic partiality. In this section, I'll suggest a line of response to exclusivity argument and the no-safe-dose argument.

First, we can concede that epistemic reasons are the only kinds of reasons that provide epistemic justification for belief, but that doesn't mean that epistemic reasons provide the only reasons for belief. But that's a slightly different issue, which I'll not discuss here.⁶³ It's the next part of the exclusivity argument that appears vulnerable. In particular, the claim that if there are only epistemic reasons to believe that p , then there are only epistemic reasons to bring it about that you believe that p is dubious. As Piller explains, '[o]ur reasons are evidential but the way we deal with these reasons is influenced, amongst other things, by our commitments' (2016, 333). It's not clear to me, for example, how the exclusivity principle would exclude us from having moral or prudential reasons, generated by our friendships, for engaging in practices that incidentally make a difference to what we believe.

Indeed, we might think that epistemic reasons to believe that p (or to bring it about that you believe that p), on the one hand, and prudential and moral reasons to believe that p (or to bring it about that you believe that p) are simply incommensurable.⁶⁴ This view has some intuitive appeal, at least. As Feldman writes:

Suppose that one belief is prudent for me—it will maximize my well being or what I care about or what is important to me—but it is not a belief I epistemically ought to have since I lack evidence for it. We can understand the idea that forming an unjustified belief might be imprudent, since it might foster a bad habit. So, we can imagine weighing the short-term prudential gain against the long term prudential cost. But I can see no values to which we could be appealing when we ask whether the overall prudential benefit trumps the epistemic cost. (Feldman 2000, 693)

⁶³ For discussion, see Pamela Heironymi 2005, Andrew Reisner 2009, and Susanna Rinard 2017; forthcoming.

⁶⁴ Thomas Kelly (2003, 619) labels the principle that epistemic reasons and moral/pragmatic reasons are rationally incommensurable the 'incommensurability thesis'.

Imagine that a certain kind of life-saving medicine has the side-effect of causing hallucinations, which will in turn lead to the formation of epistemically irrational beliefs. We plausibly have some epistemic reasons not to take the medicine, at least on the broad understanding of epistemic rationality assumed in this chapter: We will form lots of epistemically irrational beliefs and that's bad. But our *pro tanto* prudential reasons for taking that medicine could easily outweigh any epistemic costs that are incurred by taking it. So, even though if it is epistemically irrational to take the life-saving medicine, it's not at all clear that it's necessarily irrational to do so, all things considered.

We can apply the same sort of thinking to a case of epistemic partiality in friendship. It might be epistemically irrational to give your friend the benefit of the doubt, for instance, because you possess strong evidence of their culpability for some misdemeanour or other. Defenders of limited epistemic partiality in friendship can concede that. But that doesn't mean that they must also concede that your epistemic reasons not to give your friend the benefit of the doubt outweigh any other reasons you might have for doing so.

This seems to open the way for a modest response from the defenders of epistemic partiality in friendship. When the moral demands of friendship and the demands of epistemic rationality conflict, there's no straightforward answer to the question what you should do. For Keller, it really depends on contingent factors about specific persons, in particular, about what matters more to them in a given situation in which the conflict arises. He writes:

[T]here are reasons to wish for epistemic responsibility and reasons to wish for the kinds of beliefs that produce certain goods of friendship. In some situations in which the two conflict, it will make sense to favor the norms of friendship, and in others it will make sense to favor epistemic norms. (Keller 2004, 346)

How we decide whether it's more important in a given situation to hold epistemically rational beliefs or to protect our friends' well-being by treating them with epistemic partiality remains an open question. It's possible that that this is a matter of moral or prudential reasons. Sometimes, it's overall more prudential or more morally important for us to be epistemically rational. At other times, it may be that the importance of protecting and promoting a friend's well-being outweighs the costs of epistemic irrationality, whatever they may be. But however we answer this question, the fact that epistemic reasons are the only reasons that provide epistemically justified belief doesn't give us a reason to think that epistemic partiality in friendship is epistemically irrational, all things considered.

Now let's consider a response to the no-safe-dose argument. According to the no-safe-dose argument, the costs of epistemic irrationality swamp any possible value that might be gained by apparently harmless acts of epistemic partiality. Once truth is dethroned and belief, desecrated, we risk unforeseeable, negative epistemic, prudential, and moral consequences. It seems to me that the no-safe-dose argument overestimates the risk of small transgressions.⁶⁵

Let's consider a prudential analogy of the no-safe-dose argument. Sometimes people act imprudently. Some acts of imprudence turn out to be nearly harmless, where others turn out to be extremely costly. Having a second, unnecessary cup of coffee and being late to class as a result is imprudent, assuming that you wanted to get to class on time. Having a second, unnecessary cup of coffee and being late an urgent hospital appointment is considerably more imprudent, especially given the urgency of your medical problem. But it's implausible that imprudent coffee-drinking habits of the first type lead to imprudent coffee-drinking habits of the second type. A small transgression against prudence doesn't open the way for greater, more consequential transgressions.

The same seems to be true of small acts of epistemic irrationality that we find in friendship. In the final section of Chapter One, we considered the limits of Martins' epistemically partial treatment of Lime in *The Third Man*. When he thought that Lime was accused of a petty crime like smuggling cigarettes, Martins did what he could to protect his view of Lime's moral character. At first, this meant denying that Lime was involved in smuggling at all. Later, this involved changing his evaluative standards and denying that Lime was morally blameworthy for his involvement in the black market. This all changes, though, when Martins learns the true, grave nature of the allegations against Lime. In an instant, it seems, Martins ceases to give Lime the benefit of the doubt. His prior epistemic irrationality does not, it seems, prevent him from forming epistemically rational beliefs once the moral stakes are raised.

Finally, I've argued that we should treat our friends with limited epistemic partiality. Indeed, it's plausible that there's a certain wisdom to treating your friends with epistemic partiality. You need to be aware of when your beliefs and judgements could harm them and how treating them with epistemic partiality can mitigate that harm. What's more, you need to be aware of the external moral and prudential factors that count against treating your friends with

⁶⁵ James certainly has his doubts regarding the priority of epistemic rationality in our interests. Responding to Clifford, he writes 'For my own part, I have also a horror at being duped; but I can believe that worse things than being duped may happen to a man in this world: so Clifford's exhortation has to by ears a thoroughly fantastic sound' (1897/1992, 469).

epistemic partiality. But those factors don't rule out the rationality of treating your friends with epistemic partiality, rather, they impose limits on it.⁶⁶ But this brings us back to where we started. I've argued that we should treat our friends with limited epistemic partiality. What the epistemological objection shows us, if anything, is that the epistemic partiality with which we treat our friends should be limited. But the objection hasn't even shown that epistemic partiality is limited by the bounds of epistemic rationality. Rather, it shows us that when determining whether someone should treat their friend with epistemic partiality, we should take into account the comparative costs and values of epistemic (ir)rationality against the importance of protecting our friends' well-being.

§7 Summary

According to the epistemological objection, it's irrational to treat your friends with epistemic partiality, all things considered. The epistemic costs of epistemic partiality in friendship are underestimated, so the thought goes, and if they're accounted for properly, we'd see that epistemic partiality never pays off in the long run. In this chapter, I've suggested that even if epistemic partiality is epistemically irrational, it's still implausible that there that limited epistemic partiality in friendship is irrational, all things considered. What's more, I've argued that even if epistemic partiality is epistemically irrational, it's not clear that that implies that epistemic partiality is irrational all things considered. First, I argued that it's not clear that epistemic reasons and moral reasons are commensurable in the way that's required for the objection to work. If, as I suspect, they're incommensurable, there's no simple way for any epistemic reasons we might possess against epistemic partiality in friendship to override our *pro tanto* moral reasons for treating our friends with epistemic partiality. Second, I argued that the costs of epistemic irrationality, and in particular, the costs of the kind of epistemic irrationality involved in limited epistemic partiality in friendship, seem to have been overstated. Localised and non-accidental acts of epistemic irrationality don't corrupt our epistemic lives to anything like the extent that seems to be suggested. It's possible that there's a certain wisdom to epistemic partiality, I suggested. It's important to know when to give your friend the benefit of the doubt and when to hold them to account.

⁶⁶ Indeed, in the following chapter I'll consider the moral limits of epistemic partiality in friendship.

Chapter Four

A moral objection to treating friends with epistemic partiality

§1 Introduction

When Martins learns that Lime is accused of crimes that would implicate him in the deaths of many innocent people, he ceases to give Lime the benefit of the doubt. His motivation to protect his friend from the harm of unduly disparaging moral criticism is depleted almost immediately. And in *The Reader*, when Michael hears the case against Hanna, he's unmotivated to deny the allegations against her, in spite of whatever feelings they used to have for one another. Why is this? One way to see this is that he reaches the moral limits of epistemic partiality in friendship and close, personal relationships.

In the Chapter Two, I argued that we should treat our friends with limited epistemic partiality on the basis that we have *pro tanto* moral reasons to do so. In this chapter, I'll consider whether the conclusion of the argument is untenable for moral reasons.⁶⁷ As I argued in the previous chapter, your care for the well-being of your friends is a source of *pro tanto* moral reasons for you to refuse the testimony of others, specifically, when accepting their testimony would damage your friends' well-being. However, while this practice might help protect your friends' well-being, refusing someone's testimony can constitute a special kind of insult towards them (Hazlett 2017; Malcolm 2018). It seems, then, that friendship can require us to insult people by refusing their testimony, in particular, when believing their testimony would be harmful to a friend's well-being. This should count against treating your friends with epistemic partiality, or so you might think. It's plausible that we have some moral reasons not to insult other people, whether they're our friends or not.

The risk of insulting non-friends, it turns out, is just the beginning of the problems for the thesis that we should treat our friends with limited epistemic partiality. Recent discussions in

⁶⁷ The question of whether it's morally permissible to treat our friends with partiality and, if so, why, is much discussed in contemporary moral philosophy. In this thesis, I'll not attempt to add much to this debate nor to depend too much on its conclusions, except on the mainstream view that *somehow* it's morally permissible to treat friends with partiality in ordinary circumstances. For key contributions to and discussion of this debate, see Keller 2013, Kolodny 2010, Williams 1981, and Susan Wolf 1992, among others.

social epistemology, led by Miranda Fricker (2007) and José Medina (2012), among others, have brought to light the nature of distinctly *epistemic* injustice. Epistemic injustice occurs, to borrow Fricker's words, when someone is wronged 'specifically in her capacity as a knower' (2007, 20). There are several ways in which people can suffer epistemic injustice. In this chapter I'm going to focus on testimonial justice. In particular, I'll use a framework of testimonial injustice proposed by Kristie Dotson (2011), which introduces the concepts of silencing and smothering. In Chapter Two I defended the view that we should treat our friends with limited epistemic partiality. In Chapter Three, I considered and responded to an epistemological objection to epistemic partiality in friendship. Now it's time to think about the moral limits of epistemic partiality. In this chapter, then, I'll try to establish a principled way to determine when epistemic partiality in friendship is morally required and when it's morally impermissible.

The chapter will proceed as follows. First, I'll argue that treating our friends with epistemic partiality can involve insulting other people, in particular, by refusing their testimony. Then I'll consider whether we ever have *pro tanto* moral reasons to commit more morally serious, testimonial injustices in the name of friendship. After considering some putative cases, I'll introduce the concept of testimonial injustice, referring primarily to the work of Fricker (2007). I'll explain how and why the epistemic bias which I've prescribed in friendship could be thought to cause such injustice. Drawing on the work of Dotson (2011) and Sheila Lintott (2015), I'll consider some ways in which the extent of the testimonial injustice caused by bias in friendship might be worse yet, namely, because of the role that homophily plays in determining who we become (and remain) friends with. In particular, I'll consider whether epistemic partiality perpetuates racist or sexist forms of testimonial injustice. Finally, I'll attempt to alleviate some of these concerns and determine when we're morally required to treat friends with epistemic partiality. I'll argue that whilst protecting a friend's well-being might require you to cause testimonial injustice, such injustice would be morally unjustified. Accordingly, if you commit such injustice in the name of friendship, you're morally blameworthy for doing so. This, I suggest, imposes an additional limit on epistemic partiality in friendship.

§2 Can showing epistemic partiality to our friends cause us to insult non-friends?

Treating your friends with epistemic partiality can cause you to insult non-friends, in particular, when you refuse to accept what they tell you about your friends; at least, that's the claim I'll attempt defend in this section. I'll start by presenting an example, based on Stroud's (2006) Sam example. Then I'll offer some considerations in support of the claim that Sam's friends insult the

speaker by refusing her testimony, particularly drawing on the work of Hazlett (2017) and Malcolm (2018a).

So far in our discussions about Stroud's Sam example (2006), we've focussed on what Sam's friends owe him and why. We considered how it might harm Sam, if his friends don't give him the benefit of the doubt when they receive accusations against his good character. Now we'll consider the case from the perspective of the third-party speaker who reported Sam's misbehaviour to his friends.

Imagine that you've got something important to tell someone: According to your evidence, their friend, Sam, knowingly broke the heart of someone you care about after they had a one-night stand. Your intentions are sincere. You can't confront Sam directly (that would be too hostile), but you think you should tell his friends. Maybe they can explain the error of his ways to him. You go ahead and tell Sam's friends what you've heard about his behaviour, but they're not convinced. It can't be true, they say. In sum, they refuse your testimony. How does this make you feel? You might feel insulted. That's because, from your perspective, your good-faith offer of information was refused for bad reasons: How dare they doubt your credibility? In fact, not only might you feel insulted by their refusal to accept your testimony, but also their refusal really insults you.

First, we must clarify whether you really can insult someone by refusing their testimony. Under what conditions do you insult someone by refusing their testimony? You insult someone by refusing their testimony when your refusal expresses doubts about their credibility, at least according to Hazlett (2017, 42) and Malcolm (2018a). Credibility, in this context, is understood as comprising the qualities of trustworthiness and competence.⁶⁸ In the epistemology of testimony, trustworthiness is typically understood to be the quality of consistently telling people what you take to be true, or justified, or knowledge. Competence is the quality of reliably coming to form true beliefs, or justified beliefs, or knowledge.⁶⁹ This puts us in promising position to establish necessary and sufficient conditions for the special insult of refusing someone's testimony:

⁶⁸ Fricker (2007) proposes a model of credibility in these terms. Her model is very similar to one proposed by Bernard Williams (2002). Where Fricker uses the terminology of competence and trustworthiness, Williams uses the terminology of accuracy and sincerity.

⁶⁹ Hazlett (2017) and Malcolm (2018a) disagree about whether the rationality of your doubts makes a difference to whether they're insulting. Malcolm thinks so; Hazlett does not.

You insult a speaker by refusing her if and only if your refusal expresses your unreasonable doubts about the speaker's credibility.

Let's flesh this out a bit, given the view of credibility described above. You can insult a speaker by refusing their testimony. First, you insult someone if your refusal of their testimony expresses the fact that you doubt that they actually know what they take themselves to know. Refusing someone's testimony in these conditions is an attack against their competence. Second, you insult a speaker by refusing their testimony if your refusal of their testimony expresses the fact that you doubt that they're sincere. Refusing someone's testimony in this way attacks their trustworthiness. But of course, the insult of refusing someone's testimony isn't so simple. It's not the case that you insult someone every time you refuse their testimony—after all, sometimes you have very good reasons to do so. I'll now take a moment to explain Hazlett (2017) and Malcolm's (2018a) views in more detail.

To understand Hazlett and Malcolm's accounts of the special insult of refusing testimony, it's important to have a clear understanding of the relevant notion of credibility and how this fits into our understanding of testimony. Malcolm turns to Fricker (1998, 2007) for this, so I'll briefly summarise the relevant aspects of her view. According to Miranda Fricker (1998, 163), 'knowledge [...] is enshrined in the figure of the good informant'. The good informant, on Fricker's view, is someone who possesses both credibility and rational authority (Fricker 1998, 167). Credibility, as described above, comprises competence and sincerity. Competence is the property of being a reliable source of truth. Sincerity is the property of reliably sharing the truth, once you've got it. Good informants must be competent and sincere, of course, but that's not enough. Their intended audiences also need to know whether they're competent or sincere. It's no good being a credible source of valuable information if nobody recognises you as such. It's in these terms that we can start to unravel the special insult of refusing someone's testimony. First, we might refuse someone's testimony because we judge that they're not very likely to know what they take themselves to know. When we refuse someone's testimony for these reasons, what does that signal to them? Malcolm writes:

The hearer implies that the speaker lacks competence in a fundamental area of human intellectual life. She implies that he has failed to attain knowledge with respect to his belief, and since the speaker takes himself to have knowledge – to have acquired his belief in a justified way – it is an insult to treat him as though he has not. To do so is to imply that he is incompetent as a knower

with respect to this testimony – that the justification he takes his belief to have is no justification at all. (Malcolm 2018a, 54)

One reason why refusing someone's testimony is insulting is because it expresses your negative judgement about their competence. The thought at the bottom of all this is quite straightforward: In many aspects of life, it's insulting to be deemed incompetent when, in fact, you're competent. Our epistemic lives are no exception, or so the thought goes. Imagine that a tourist visiting your hometown refuses your testimony when you give them directions to their hotel, because they doubt your sense of direction. If their refusal of your testimony expresses an unreasonable judgement that you're incompetent, then their refusal of your testimony expresses an insult. In short, according to Malcolm's view, their refusal insults you. Of course, it's not obvious that we always insult people when we refuse their testimony, even if we reject it because we judge them to be incompetent. Possibly, it's not an insult to refuse someone's testimony for this reason if they really are incompetent, or your judgement that they are is reasonable. Indeed, this consideration is central to Hazlett's view of the special insult of refusing someone's testimony (2017). I'll return to consider the plausibility of this view later.

Sometimes, though, when you refuse someone's testimony, it's not because you doubt their competence. Indeed, you might think that they're very likely to know what they're talking about, but nevertheless refuse their testimony. Rather, you might refuse someone's testimony because you think that they're not trustworthy. In situations like this, you judge that you can't rely on the speaker to tell you the truth, even though you suspect it's in their possession. Correspondingly, your refusal of someone's testimony insults them if it expresses an unreasonably negative judgement about their trustworthiness. Malcolm writes:

If a hearer treats a speaker as though he is attempting to intentionally deceive her – as though he is lying with his testimony – she may insult him by implying that he is performing a morally objectionable action. As with other immoral acts, it is offensive to accuse someone of performing one. So, it would be offensive to suggest that someone is a thief, or abusive, or racist, or a liar, even if it turns out that she is none of these things. As such, when we reject a speaker's testimony, and do so because we take the speaker to be lying, we can insult the speaker by implying that she has acted immorally. (2018a, 56)

Here we have two, compatible explanations of how you can insult someone by refusing their testimony. Of course, as Malcolm points out, the reason you refuse their testimony are relevant. It can't always be the case that refusing someone's testimony is an insult.

There are many ways we can reasonably come to doubt someone's credibility. Malcolm (2018a) identifies three primary sources of reasons for refusing someone's testimony, namely, defeaters, inductive reasons, and prejudgement reasons. I'll begin by explaining Malcolm's view that defeaters can give you a reason to doubt someone's testimony and what this means for our understanding of the insult of refusing someone's testimony (2018a, 57-58). Sometimes you're just in a better evidential position than the speaker. Their evidence for believing that p is outweighed by your evidence to the contrary. It's epistemically rational to refuse their testimony in these conditions. If you have strong evidence for believing that not- p , then, when someone with weak evidence for believing that p tells you that p , you should refuse their testimony. In this situation, you possess rebutting defeaters for their testimony. Undermining defeaters are different. Rather than weighing against your interlocutor's evidence, undermining defeaters weaken the evidential-support relation between their evidence and their belief. Whilst the speaker believes according to their evidence, you're in possession of additional information, according to which their evidence offers diminished support their belief. When you possess defeaters of either kind for a speaker's testimony that p which they don't possess, it's not so obvious that you insult them by refusing their testimony. Certainly, it would seem unreasonable on the part of the speaker if they *felt* offended by your refusal of their testimony, at least once they'd been informed that you possess overwhelming defeaters for the justification of their belief.⁷⁰

There are occasions, as Malcolm points out, in which you refuse someone's testimony because of the defeaters you possess but your refusal nevertheless constitutes an insult towards them. He offers the following example:

Imagine a friend who sides with a political party whose leader is a disingenuous and shift character. Even though you are both aware of this, and as such you have a defeater for the leader's testimony, your friend is taken in by the things he says, gullibly believing many of the claims he makes to support his policies. Now your friend is reeling off "facts" and "statistics" to you, telling you to believe her testimony about them so that you will vote differently. You refuse

⁷⁰ That said, we should be careful to distinguish between the question whether someone feels insulted—whether they take offense, as it were—and whether your actions constitute an insult. If you viciously mock someone who has no self-respect, they may not feel insulted by your words. Nevertheless, your vicious mockery is insulting. As Malcolm points out (2018a, 56), the same is true when someone insults you in a language you don't understand: You're insulted whether or not you understand the meaning of their words or their intention in uttering them.

to believe her, insisting that you think she's been deceived. Your reasoning behind this is that she's got her judgments all wrong on this politician, and he's clearly just fabricating the facts for political gain. You imply that she's failed to distinguish truth from lie, perhaps under the influence of persuasion and manipulation, and has displayed culpable failure of epistemic competence. Since she has a defeater for her beliefs – the evident dishonesty of the politician – you judge that your friend has self- deceptively put this aside to the point of irrationality, and this sounds offensive to some degree. (2018a, 58)

On Malcolm's reading of this case, you insult your friend by refusing her testimony. You insult the speaker in this case because your judgement that they've failed to respond reasonably to a defeater that you both possess expresses an unreasonable doubt in the speaker's competence. To the extent that the failure to respond reasonably to a defeater that you possess (for instance, by suspending your belief) demonstrates intellectual incompetence. To refuse someone's testimony because you unreasonably judge them to have demonstrated this incompetence is, therefore, an insult to them.

Returning to the example that has been central to this chapter, we can ask what kinds of defeaters is it plausible that Sam's friends might have for the third party's testimony about Sam? It's plausible that Sam's friends would have rebutting defeaters for her testimony. Indeed, it's tempting to think that as his friends, they might be better-positioned to ascertain the truth about their friend Sam and his moral character than the third-party speaker. If, according to their evidence, it's reasonable for them to believe that Sam is a caring person, then this will go some way to rebut the allegations of the third party. We can imagine how this would play out in practice. Sam's friends' reasoning could go something like this:

- (1) If the third party's testimony is correct, then Sam's not a caring person.
- (2) But Sam is a caring person, and we should know; we're his friends.
- (3) Therefore, the third party's testimony is not correct.

For Sam's friends there are two possible explanations of the speaker's utterance. Either the speaker is incompetent or she is insincere. With sufficiently strong prior evidence that Sam is a caring person, his friends are in a strong evidential position to refuse the allegations of the third party.

But this way of rebutting criticism of our friends won't always work. The thing is, we can't speak for every case, or even the majority of cases, simply by pointing out that it's possible that Sam's friends could have legitimate rebutting defeaters for the speaker's testimony. Indeed,

granted a pessimistic view of human nature, we might think that it's quite improbable that Sam's friends have legitimate rebutting defeaters for the testimony of the third party. As Stroud herself writes (2006, 516), some of the force of the example of Sam's misdemeanour come from the fact that people are quite prone to selfish or cruel behaviour in their sex-lives. Indeed, in Chapter One I described a range of epistemic behaviours that would explain why Sam's friends have an unreasonably inflated sense of his good moral character. This is relevant because it suggests that we shouldn't be so sure that Sam's friends will have strong enough evidence for Sam's good moral character to rebut the third party's testimony, at least, not in the majority of cases. So, when Sam's friends think they possess defeaters for the testimony of the third party, it's plausible that their confidence in the strength of those defeaters is inflated. Now, if that's the case, and they refuse the speaker's testimony on that basis, then their refusal of the third party's testimony is an insult to her.

Of course, it might be the case that you refuse someone's testimony and thereby insult them, but not because you're unreasonably confident that you possess defeaters for their testimony. Sometimes, we refuse a speaker's testimony because we have doubts about their credibility, owing to the fact that they've deceived or misled us in the past. With this consideration in mind, I'll now explain why Malcolm (2018a, 58-60) thinks that inductive reasons can give you a grounds to doubt someone's testimony. Then I'll consider what this means for our understanding of the insult of refusing someone's testimony.

If a speaker has demonstrated incompetence in the past, then, plausibly, it's epistemically rational to doubt their credibility. Likewise, if a speaker has demonstrated their untrustworthiness in the past, then, plausibly, it's also epistemically rational to doubt their credibility. Plausibly, then, what you believe about someone's track record can give you good epistemic reasons to doubt their credibility. In reality, of course, things are a little more complicated. For one thing, it's worth noting that you can at a given time have inductive reasons for doubting someone's credibility without their having a bad track record. A speaker's credibility is temporarily compromised if, for instance, they're drunk, or they're being blackmailed. The thought is that if you possess inductive reasons to doubt someone's credibility, then your refusal of their testimony doesn't insult them so strongly (on Malcolm's view) or at all (on Hazlett's view). For example, if you have good evidence that a speaker is an incompetent arithmetician, then you shouldn't accept their testimony when they tell you how much you owe after you decide to split the bill for dinner. Likewise, Malcolm points out (2018a, 59), this is why it's not such a strong insult to refuse someone's testimony in boy-who-cried-wolf cases. If you've got a

good reason to suspect that someone's lying to you, then it's not such a strong insult to be wary of accepting what they tell you.⁷¹

Now we should consider how this way of insulting people by refusing their testimony fits into the wider project of this chapter. Is it plausible that Sam's friends have inductive reasons to refuse the speaker's testimony? And if so, when, if ever, does their response to these reasons insult the third-party speaker? Let's keep working with Stroud's Sam example (2006). Suppose that Sam's friends unreasonably believe that the speaker lacks credibility. They don't have sufficient inductive reasons to support their judgement. Under these conditions, if they refuse the speaker's testimony, then they insult them. Given the theory of testimonial insult presented here, they insult them because their refusal of their testimony expresses an unreasonably critical view of their competence or trustworthiness.

Unfortunately, this doesn't look like the most promising way of understanding the insult of refusing someone's testimony for reasons of friendship. We can imagine that before hearing the third party's judgement about Sam's behaviour, his friends have a reasonable, positive judgement about the third party's credibility. For them to insult the speaker by refusing their testimony about Sam's behaviour for inductive reasons, they would have to have reached an unreasonable, negative judgement about their testimony on the basis of inductive reasons before she tells them what Sam did. Otherwise, it's hard to see how inductive reasons play a part in their refusal of their testimony at all.

Finally, I'll explain how prejudgement reasons lead us to doubt a speaker's testimony. I'll argue that this account of the insult of refusing someone's testimony best fits the cases we've discussed so far in this chapter. Moreover, I'll argue that it best explains the nature of the insult caused when someone's testimony is refused for reasons of friendship.

We often seem to rely on prejudgement reasons in our reasoning about other people's testimony. In fact, it's common enough that we only have prejudgement reasons to inform our reasoning. In support of this claim, Malcolm (2018a, 60) quotes Fricker:

Barring a wealth of personal knowledge of the speaker as an individual . . . [a hearer's] judgement of [a speaker's] credibility must reflect some kind of social generalization about the epistemic trustworthiness – the competence and sincerity – of people of the speaker's social type, so that it is inevitable (and desirable) that the hearer should spontaneously avail himself of the relevant generalizations in the shorthand form of (reliable) stereotypes. (Fricker 2007, 32)

⁷¹ That said, it's a mild insult, one we might call a slight, according to Malcolm (2018a, 51).

To refuse someone's testimony for prejudgement reasons is to refuse someone's testimony because of judgements you have about their credibility which are based on your understanding of their social identity. Malcolm carefully uses the term 'prejudgement' rather than prejudicial to emphasise the fact that not all cases in which a hearer has reasons to doubt someone's testimony based on their social identity are cases of deep injustice (2018a, 60). The speaker's social identity might comprise an intersection of factors, including their sex, their race, and their religion. I would suggest that we consider your friendships as factors which contribute to your social identity too. After all, along with sex, race, and religion, your friendships are part of your social life; they can and are tracked by other people as a means of making judgements about you.

But what can this account of the insult of refusing someone's testimony tell us about the cases in which a speaker's testimony is refused by the listener in order to protect the listener's friend? Well, when Sam's friends refuse the testimony of the third party, it's plausible that they do so for prejudgement reasons. The speaker's testimony is prejudged and, ultimately, refused because of her social position. The relevant aspect of the speaker's social position is that she is a non-friend. When Sam's friends identify that the speaker isn't one of Sam's friends, they judge that they are therefore not a credible source of some information about Sam. Indeed, the deficit in the credibility they attribute to the speaker is likely to be subject-specific: They pre-judge that non-friends are not credible sources of moral criticism of friends.

When the third party tells Sam's friends that he slept with someone and knowingly broke their heart, his friends refuse their testimony. More specifically, it's because they're not Sam's (or their) friend that they don't accept criticism of Sam from them. We might imagine how they rationalise this through the mechanisms of biased interpretation that were discussed in Chapters One and Two: 'Since they're not Sam's friend, they simply aren't in a position to know', they might say. What's more, they might reason that it's neither her business to judge Sam's behaviour ('They're not even his friend!') nor to share their judgements with them ('And they're not our friend either!'). Indeed, armed with these assumptions about who does and who doesn't have the right to form an opinion about Sam's character and conduct, his friends might construe the very fact that they, a non-friend, have shared their information about Sam's conduct as a reason not to believe them.

How, then, do Sam's friends insult the third-party speaker when they refuse their testimony? We can suppose that the strategies of epistemically partial, evidence processing that Sam's friends use to rationalise their prejudgement of her criticism are epistemically defective. Accordingly, the judgements they end up with are epistemically unreasonable. Suppose, as before, that the third party reasonably believes that Sam slept with someone and then knowingly

broke her heart. What's more, suppose that the third party is quite sincere in their effort to tell Sam's friends about his behaviour. The third party, then, deserves credibility, but they're not treated as a rational authority, despite the fact that they deserve to be respected as such. Their unreasonable refusal of their testimony is an insult.

It seems, then, that friendship can be a source of *pro tanto* moral reasons to insult non-friends by refusing their testimony, in particular, when the alternative, that is, accepting their testimony, would damage either our friend's well-being. In the section that follows, I'll consider whether treating friends with epistemic partiality risks causing testimonial injustices. I'll argue that it does, drawing on the work of Kristie Dotson (2011). Against this theoretical background, I'll consider and respond to an objection owed to Sheila Lintott (2015), who claims that epistemic partiality in friendship could be morally problematic for reasons because it risks perpetuating injustice.

§3.1 Can showing our friends epistemic partiality perpetuate injustice?

Refusing to accept someone's testimony can constitute a harm more serious than insult, particularly when the content of that testimony relates to serious moral misconduct.⁷² Let's start with an example. Imagine that, in a private conversation, a non-friend (hereafter the speaker) accuses your friend of serious moral misconduct. Suppose that your friend is accused of repeatedly using a racial slur towards a third party. This is, of course, a morally grave version of Stroud's Sam case. Let's suppose, as before, that the speaker's belief about Sam's conduct is reasonable and that she has no reason to make false accusations. As in Chapters One, Two and Three, we can start by supposing that this is one of those cases in which it's reasonable at least for a neutral hearer to accept the speaker's testimony. The hearer doesn't have sufficient non-testimonial evidence to the effect that the speaker is unreliable or trustworthy. The point is this: By all accounts of testimony, a neutral party should accept the accusation.

But what should a friend do when presented with such a serious accusation of moral misconduct? According to the argument presented in Chapter Two, we should treat our friends

⁷² The idea that the demands of a friendship could conflict with the demands of morality are not unprecedented. Dean Cocking and Jeanette Kennett (2000) argue that friendships can involve an element of 'moral danger', and that friendships can be valuable in spite of this fact. Relatedly, Daniel Kolstonski (2016) argues that being a good friend can require you to help a friend in committing morally objectionable acts.

with epistemic partiality. In the case discussed above, borrowed from Stroud (2006), the harm that would be caused by incorrectly believing the accusation against your friend outweighs the harm that would be caused by incorrectly rejecting the conclusion. Now it's plausible that if you come to believe falsely that your friend repeatedly used a racial slur to abuse someone, you harm your friend. According to the broad account of how beliefs harm, it's likely that it would harm your friend if you believe a false accusation of serious moral misconduct. As it happens, though it's not a vital premise of the argument here, it's possible, given the broad account of beliefs that harm from Chapter Two, that believing the accusation would harm your friend even if the accusation is true. Your friend may desire that you're loyal to them or that you think well of him, and believing the accusation would frustrate those desires.

It's also plausible that if you believed that your friend had racially abused a third party, then your commitment to friendship with them would be impeded. That's because, plausibly, learning that your friend is racist will cause you to feel what Jones (2011) calls lover's shame. Coming to feel ashamed of someone can cause you to withdraw your commitment to the relationship which you share with them, thereby frustrating your friendship with them. To put it simply, what this tells us that as far as protecting your friendship with the accused is your priority, it's likely you should reject the accusation.⁷³

It hardly needs stating that this conclusion is morally objectionable. It would be an unacceptable implication of the argument from care presented in Chapter Two, that we ought to ignore accusations of serious moral misconduct against our friends. But how do we make sense of this, now that we're already committed to the claim that beliefs (and omissions of belief) can harm? It's worth considering the specific way in which refusing the testimony of a genuine victim can be harmful.

Why, then, is it so dangerous to refuse this kind of testimony? First, let's work within the framework established in the first half of this chapter, according to which you insult someone by refusing their testimony if and only if your refusal expresses unreasonable doubts about their credibility. When the speaker shares her testimony with someone in this case, it's highly unlikely that the hearer will possess defeater-reasons, inductive reasons, or prejudgement reasons, to

⁷³ I say 'likely' because, as discussed in Chapter Two, there are circumstances in which believing something negative about your friend won't harm them, or the harm will be an unavoidable consequence in pursuit of the 'greater good' for your friendship. A friend who genuinely seeks to reform their behaviour might not be harmed in the same way that a friend who denies there's anything wrong with their racist attitudes.

refuse her testimony. Since it's unlikely that the hearer possesses good reasons to refuse the testimony of the speaker, it's likely that the hearer insults the speaker by refusing her testimony. It's worth considering the severity of the insult paid by the hearer in these circumstances. Arguably, it's extremely insulting to refuse someone's testimony in a case like this because of the gravity of the implication made by that refusal: It implies that the hearer judges that the speaker is either not competent enough to know what happened to her, or otherwise is insincere. It's a very serious insult to judge unreasonably that a victim of racist abuse is lying about what happened to them. The insult is so much more serious because the judgement about their credibility is so much more critical. To doubt that someone is telling the truth about an incident of racist abuse is a serious insult because of how bad it would be to lie about such an event. To deny that someone is competent enough to know whether someone else has been the victim of racist abuse is a serious insult, at least because of the high level of incompetence implied.

We can also consider the negative social and political consequences of refusing the third-person reports of racially or sexually aggravated violence or discrimination. It's highly plausible that rejecting reports of racism can be extremely harmful, not only to the victim herself, but also to other victims. By refusing to accept a report of racism, you immediately deny the victim a vital opportunity for support. It's also reported that many victims of racism are reluctant to report the abuse they have suffered, because they believe that their reports will not be accepted and that, if they are accepted, they will be misunderstood. (We'll return to this topic in the next section.) The point is that even before we've delved very deeply into the ethics of bias, we should be reluctant to accept a theory of epistemic bias in friendship which would promote rejecting such testimony.

What I've tried to show in the previous section is that refusing morally important testimony silences voices that ought to be heard. It would be extremely problematic if friendship required this. The thing is, to the extent that your care for your friend is frustrated by your forming negative moral judgements about them, friendship does appear to require you to silence these voices. In the following section, I'll explain in more detail how this fits into the theoretical framework of testimonial injustice.

§3.2 Testimonial silencing

In the previous section I considered a case in which morally serious testimony was, in a sense to be explained now, silenced. The silencing of this testimony appears to be a result of the demand for epistemic partiality in friendship. In the example, Sam's friend is told that Sam used a racist slur. Because of the hearer's friendship with Sam, and in particular because of their concern for

Sam's well-being, the speaker's testimony is undervalued; she's not given the credibility that she merits. We considered the possibility that the speaker would withhold her testimony as a result of her believing that it wouldn't be accepted anyway. Both of these cases seem like cases of serious injustice. In the section that follows, I'll explain why it's correct to understand these as cases of a distinctly epistemic kind of injustice, namely, testimonial injustice.

What's going on here, exactly? Dotson (2011) identifies two practices of silencing, namely, quieting and smothering. I'll explain each of these practices of silencing and then explain why, drawing on Lintott's (2015) work, we might worry that a demand for epistemic partiality in friendship constitutes a demand for testimonial injustice. A speaker is quieted when her testimony is undervalued or even completely ignored by its audience. When a person has been quieted, her audience doesn't recognise her as a good informant. Dotson writes,

The problem of testimonial quieting occurs when an audience fails to identify a speaker as a knower. A speaker needs an audience to identify, or at least recognize, her as a knower in order to offer testimony. (2011, 242)

Among the principle causes of testimonial quieting is the phenomenon of negative epistemic stereotyping. This is, according to Fricker (2007), a form of identity prejudice. Consider how stereotypes undermine the testimony of women on matters which supposedly require intellectual rigor and cool-headedness.

Testimonial injustice, according to Fricker, arises when someone's testimonial contribution is undervalued as a result of identity prejudice. For Fricker,

A negative identity-prejudicial stereotype is ... [a] widely held disparaging association between a social group and one or more attributes, where this association embodies a generalization that displays some (typically, epistemically culpable) resistance to counter-evidence owing to an ethically bad affective investment. (2007, 35)

Fricker proposes a model of conversational interaction which gives a foothold in our effort to understand the nature of the phenomenon of testimonial injustice. In an ordinary (and, more to the point, successful) conversational interaction, the speaker and the hearer rely on stereotypes and heuristics to make judgements about the sincerity and reliability of their interlocuter. She remarks:

This model of the interaction between speaker and hearer helps us to see the mechanism whereby identity prejudice can distort a hearer's credibility judgement: it distorts the hearer's perception of the speaker. Applying the perceptual idiom to our chief example, we can say that the judgement of the jurors of Maycomb County is so distorted by prejudicial racial stereotype that they cannot, in that courtroom context, perceive Tom Robinson as anything but a lying Negro. Now in this example the jurors' perceptions are shaped *inter alia* by prejudiced beliefs; the prejudicial racial stereotype determining their credibility judgements is in part doxastically mediated. (Fricker 2007, 36)

In the case of Tom Robinson, his testimony was received but refused. His attempt to confer knowledge failed because his audience didn't afford him sufficient credibility. Sometimes, however, a speaker doesn't get as far as sharing their knowledge at all, but rather, they withhold it. We withhold testimony for all sorts of reasons, for instance, to avoid offending someone or even simply to cut short a conversation that's taking too long to wind up on its own. Not all instances in which a speaker withholds her testimony are so innocuous. We'll follow Dotson in calling the phenomenon of the coerced withholding of testimony 'smothering'. And what is smothering? A speaker's testimony is smothered when she is illegitimately coerced into withholding her testimony. Dotson puts it thus:

Testimonial smothering, ultimately, is the truncating of one's own testimony in order to insure that the testimony contains only content for which one's audience demonstrates testimonial competence. (2011, 244)

In an archetypal case, someone who is the victim of smothering possesses the relevant intellectual virtues to justify (in some sense) their telling their audience that *p* (i.e., their belief that *p* is true and they are sincere), but they withhold their testimony because they are confident that their testimony will be refused or misunderstood anyway.

For an instance of withheld testimony to count as testimonial smothering, according to Dotson's account, three conditions must be met. The purpose of these conditions is to rule out potential counter-examples. She writes,

The three circumstances are: 1) the content of the testimony must be unsafe and risky; 2) the audience must demonstrate testimonial incompetence with respect to the content of the testimony to the speaker; and 3) testimonial incompetence must follow from, or appear to follow from, pernicious ignorance. (2011, 244)

It's worth taking a moment to consider each of these conditions. On Dotson's view, unsafe testimony 'is testimony that an audience can easily fail to find fully intelligible' (2011, 244). Risky testimony, for Dotson, 'runs the risk of leading to the formation of false beliefs that can cause social, political, and/or material harm' (2011, 244). The thought is that a piece of testimony is unsafe and risky for a given hearer if (i) the hearer could easily misunderstand the testimony, thereby coming to hold false beliefs, and (ii) if the hearer misunderstands it, then the resulting false beliefs are likely to have serious consequences. As an example of unsafe and risky testimony, Dotson (2011, 244-5) gives the example of testimony about domestic violence in non-white communities. Referring to the work of Kimberlé Crenshaw (1991), Dotson explains how, when women of colour consider speaking about domestic violence in African American communities, they often withhold their testimony because of the risk that what they say will be taken to justify harmful stereotypes about African Americans:

Some, though certainly not all, African Americans have considered the ramifications of testimony about certain kinds of occurrences, like domestic violence and/or rape, to be a detriment to African American communities at large, often at the expense of those who suffer from domestic violence and/or rape.⁷⁴ (2011, 245)

The thought is that the content of the testimony is unsafe because audiences are likely to misunderstand it, incorrectly taking reports of discrete instances of wrongdoing as evidence of more general behavioural trends. The content of the testimony is risky because, if the hearers of the testimony misunderstand it, they're likely to form damaging false beliefs about African American men.

The second necessary condition of testimonial smothering, on Dotson's account (2011, 245), is that the hearer fails to show the speaker that she is a competent recipient of testimony. Testimonial competence involves some degree of proficiency in at least the following two skills: On the one hand, your testimonial competence depends on your ability to understand what you're told, and on the other hand it depends on your ability to recognise when you don't (or you're not likely to) understand what you're told. Dotson (2011, 245) reflects on what it's like to listen to a lecture on theoretical physics as a non-expert. The competent, non-expert can not only understand some of what she's told in the lecture, but also recognise when she doesn't understand. Testimonial smothering requires that the speaker withholds or truncates her testimony because of the hearer's failure to demonstrate that she's got what it takes to interpret

⁷⁴ Dotson cites Crenshaw (1991, 1256-1257) in support of this point.

the speaker's testimony and to appreciate her own limitations. By including this condition, Dotson anticipates potential counter-examples in which someone withholds their risky and unsafe testimony because they have unreasonable doubts about the competence of the hearer. The thought is that testimonial injustice doesn't occur when someone withholds their testimony from someone who has nevertheless demonstrated that they're a competent recipient of testimony. When testimony is unjustly smothered, it will be the case that the hearer has failed to demonstrate to the speaker that she's able to understand the speaker's testimony, or recognise when she's failed to understand the speaker's testimony.

A third and final, necessary condition for testimonial smothering is that the testimonial incompetence of the hearer must result from (or appear to follow from) pernicious, situated ignorance. What, then, is situated ignorance? Someone in a state of situated ignorance lacks knowledge as a result of their social positioning.

Situated ignorance, which follows from one's social positioning, is a result of epistemic limitation that fosters a kind of epistemic distance between those not in possession of that limitation and those who do possess the limitation. (Dotson 2011, 248)

Epistemic distance between two or more people exists when there is a gap between their respective worldviews, such that they have different ways of seeing and understanding the world. Differences in race, gender, and social and economic status can all contribute to the growth of epistemic distance between persons or peoples, at least according to Dotson (2011, 248). According to Dotson's account, testimonial smothering occurs only if the failure of the hearer to demonstrate testimonial competence results from situated ignorance.

This quick summary is sufficient to demonstrate what Dotson means by the practices of silencing. Testimony can be quieted or smothered, depending on whether the hearer fails to afford the speaker their due credibility, or the speaker withholds her own testimony because she reasonably doubts that she would be afforded her due credibility. Put this way, smothering occurs as a result of anticipated quieting. The speaker withholds some or all of her testimony because she believes that testifying will not serve its intended purpose. In anticipation of the predicted failure of her attempt to convey her knowledge, the speaker refrains from speaking at all.

Let's consider how this relates to the cases we considered earlier. When the speaker tells Sam's friend that Sam used a racial slur against them, the friend refused their testimony. Unless Sam's friend has good epistemic reasons to think that the speaker is lying or mistaken, their refusal of this testimony has the appearance of an instance of testimonial quieting. Similarly, in

cases where the speaker withholds their testimony from Sam's friends because they're confident that, as his friends, they won't accept their testimony, we find the appearance of testimonial smothering.

A piece of the puzzle is missing, however, if the epistemic bias showed by Sam's friends is to count as testimonial injustice. Testimonial quieting and testimonial smothering both require pervasive ignorance or identity prejudice, but it isn't immediately obvious that that's what's happening when we refuse someone's testimony for reasons of friendship. To be clear, it isn't obvious that it's Sam's friend's racist prejudices that cause them to refuse the testimony of Sam's accuser when they say that Sam used a racial slur. Nor is it clear that it's Sam's friend's misogynistic worldview that causes them to refuse the testimony of the third party who accuses Sam of sexual assault. Although racism and misogyny undoubtedly lead to the silencing of victims, in the cases we've considered, another explanation for the refusal or withholding of testimony is available, namely, that Sam is the hearer's friend. In the section that follows, I'll make the case that, in fact, pervasive ignorance can account for the fact that we refuse testimony for such reasons. Then, I'll conclude, it's reasonable to think that there's a serious moral problem with the conclusion that we should treat our friends with epistemic bias.

§3.3 Sheila Lintott on homophily, bias, and friendship

For friendship to present a serious risk of testimonial quieting or smothering, it must be the case that something about the nature of friendship increases the risk of pernicious ignorance. In fact, as Sheila Lintott explains, it's surprisingly likely that friendship could perpetuate pernicious ignorance. To begin, Lintott (2015) draws attention to the fact of homophily in social relations. She writes:

The basic idea is that if I see another as a member of "my group," whatever group that might be, I am likely to be positively biased toward and attracted to that person. Positive emotions are more strongly felt and commitment and loyalty are more active for "in-group" people. As a corollary, I am more or less indifferent to and sometimes negatively disposed toward those I perceive as not members of my group, as, in other words, members of an "out-group." (2015, 321)

What Lintott describes here fits well with what was said in Chapter Two about nature of commitment in friendship. Reflecting on Telfer's work on friendship (1970), we discussed how feeling a shared bond is essential for friendships to flourish. To be friends with someone requires

more than just liking them; you must feel connected by some common appreciation of what's good and bad, what's right and wrong, and so on. The thing is, the conditions which create these bonds are liable to perpetuate prejudice based on social identity. Lintott continues:

Shared past experiences and interests definitely bond us: they can be shortcuts to friendship. But there is a downside to relying too heavily on innocent and obvious seeming connections, for they can be barely veiled shorthand for racial, ethnic, class, gender, or other similarities that may trigger various biases and prejudices. (2015, 324)

It's plausible that a person's social identity influences the friendships they form. If that's correct, then, at least in many cases, it seems totally innocuous. Some groups of friends bond over a shared hobby. Their engagement in a shared activity provides that 'sense of a bond' which was important in Telfer's account. This is what brings together friends who play sport together or attend the same book club. Lintott doesn't pull her punches:

[W]e need to pause to ask what determines to whom we "already feel a special concern." "Because he is my friend" may provide some moral justification for differential treatment while attempts to justify differential treatment "because he's white," "because she's Catholic," or "because he's heterosexual," patently fail. However, if a major reason a friendship has formed concerns social identity, then justifying partiality "because he is my friend" is not entirely distinct from justifying partiality "because he's white." (2015, 325)

What Lintott says about homophily and bias in friendship suggests that it's not unreasonable after all to think that the preference we have for our friends' interests over the interests of non-friends could lead to morally serious cases of testimonial injustice. Now I'll suggest an explanation why.

Why, we might ask, would Sam's friends refuse testimony which alleges that he racially abused a third party? It's tempting to answer that they would refuse the testimony because Sam is their friend and leave the matter there. However, to echo Lintott's words, we should ask ourselves to what extent the fact that Sam is their friend is explained by relevant facts about social identity, in this case, race. If the fact that Sam is the hearer's friend is caused by the fact that Sam is white, then the quieting of the third party's testimony might well be explained by social identity.

In the previous section, I considered two examples of testimonial silencing, one each of quieting and smothering, which I take to be both plausible and morally problematic. More to the

point, the examples suggest that the requirement for epistemic partiality in friendship is morally problematic because it seems to demand that we propagate testimonial injustice in the name of friendship. Drawing on the work of Dotson (2011) and Lintott (2015), I've tried to provide a framework for understanding how the practices of epistemic partiality can be seen as causes of testimonial injustice. To the extent that your friendship with someone, and specifically the fact that you're close to them, is a result of your own prejudices and/or morally pernicious ignorance, epistemic partiality in friendship risks perpetuating testimonial injustice. This is morally unacceptable, so we should carefully consider the extent to which epistemic partiality is morally permissible in friendship.

§3.4 The problem: It's immoral to treat your friends with epistemic partiality

We find ourselves in a bind. On the one hand, it seems like the harm caused by insulting someone by refusing their testimony is morally excusable in some cases, in particular, when the harm done to the speaker is minimal. On the other hand, when your refusal to accept a speaker's testimony for reasons of friendship causes epistemic injustice through testimonial silencing, it seems the harm is by no means morally excusable. Given this tension between our intuitions, it's necessary to develop a principled way of determining why it's morally permissible to pay minor insults for reasons of friendship, but not to cause testimonial injustice.

§4 The response: It's a matter of moral priorities

In this section, I'll propose a way of determining when it's morally permissible to treat your friends with epistemic partiality. I'll argue that it's a matter of moral priorities. Treating your friends with epistemic partiality is morally wrong when they stand accused of serious moral misconduct because to do so is to get your moral priorities wrong. This moral wrongness, I'll argue, is a source of *pro tanto* moral reasons against treating your friends with epistemic partiality when they stand accused of serious misconduct. However, treating your friends with epistemic partiality isn't morally wrong when they stand accused of minor moral misdemeanours.

There seems to be something morally objectionable about treating close friends with epistemic partiality when they're accused of serious moral misconduct. For instance, it's plausible that if your friend is accused of racist violence but you treat them with epistemic partiality, and thereby excuse them from blame, at least in your view, you've done something morally objectionable. I've already suggested that one reason why this is wrong is that it would

perpetuate the injustice of testimonial silencing. But it doesn't always seem morally objectionable to treat close friends with epistemic partiality. Sometimes, it's true that treating our friends with epistemic partiality risks insulting, but to the extent that this insult harms them, it nevertheless seems excusable.

This is puzzling. How do we explain the wrongness of treating friends with epistemic partiality? It's a matter of moral priorities, or so I'll argue in the following section. First, let's consider a different but closely related question: Why is it wrong to be friends with morally bad people? Jessica Isserow (2018) defends a moral-priorities account of the wrongness of friendship with morally bad people, which can be fruitfully applied to the problem under discussion in this chapter.⁷⁵ The central claim of Isserow's view is the following: While it's not always immoral to overlook a friend's flaws, it is when your overlooking their flaws causes you to lose sight of your moral priorities or renege on your moral commitments. Isserow writes:

One could understand an individual who was willing to forgive a friend's failure to recycle; for this is a fault in spite of which we could plausibly accept someone. But an individual who discounted a friend's rampant racism would suggest to us that she could not care less about the values which tell against racism, or for the potential victims of racist attitudes. At the very least, she would suggest to us that she does not stand for (or is not standing up for) such values in the fullest sense. Her willingness to discount vices of this extreme sort would suggest that there are certain values to which she is not properly responsive. (2018, 3112)

On Isserow's view, someone who is knowingly maintains their friendship with a morally bad person is improperly responsive to important moral values. It's here, Isserow claims, that they exhibit the vice of moral complacency. Moral complacency, to be clear, is the moral vice of having your moral priorities out of order. Morally complacent friends excuse what shouldn't be excused (Isserow 2018, 3112). Someone who remains friends with a known racist has their moral

⁷⁵ Isserow (2018) describes three common-sense explanations, before defending her own, more sophisticated, account. According to desert views, it's wrong to be friends with morally bad persons because they don't deserve to have the good of friendship bestowed upon them. According to abetting views, it's wrong to be friend with a morally bad person because to do so is to aid and abet their misconduct. And according to the risk view, it's imprudent to be friends with morally bad persons because to be friends with such a person is to put yourself in harm's way. Since it's not the purpose of my chapter, I'll not explain the advantages and disadvantages of these views here. For a thorough exploration of the topic, see Isserow 2018.

priorities out of order because, Isserow claims (2018, 3112), they judge the value of their friend's redeeming features to outweigh the disvalue of their racism, found in the harmfulness of their very attitudes and in harmful actions that result from their attitudes.

This puts us in a better position to understand when it's immoral to treat a friend with epistemic partiality. You exhibit moral complacency when you treat a friend with epistemic partiality with respect to an accusation of serious misconduct, such as those discussed in section 3 of this chapter. What's wrong with epistemic partiality in these cases is that you're incorrectly allowing the value of your friendship to outweigh the disvalue of their racism. In other words, you're giving your friendship too high a place in your moral priorities. The thought is that if you prioritise clearing their name rather than finding out the truth, you express a vicious lack of concern for the wrongness of racism. When it the question whether your friend is racist becomes salient, your priority should focus on establishing whether they really are racist, and not protecting their well-being.

Treating your friends with epistemic partiality when presented with accusations of serious moral misconduct is morally complacent and risks perpetuating testimonial injustice via silencing. This must be taken into consideration when considering whether we should treat friends with epistemic partiality. Now, if we take moral reasons to be a species of pragmatic reasons, then, when it would be morally complacent to treat a friend with epistemic partiality, you have moral, pragmatic reasons not to treat your friend with epistemic partiality when they're accused of serious moral misconduct. In practice, of course, this will likely turn out to be more easily said than done. You might harm your friend by refusing to give them the benefit of the doubt when they face such morally serious accusations. For instance, they may be wounded by your lack of loyalty. However, when the question whether they're racist becomes salient, it's reasonable to conclude that there's more to be lost than there is to be gained from this friendship. In summary, it's possible that the moral considerations against treating a friend with epistemic partiality outweigh the considerations in favour of it. Under these conditions, you should not treat your friend with epistemic partiality.

We might question the gravity of moral complacency. Isserow claims we don't typically rank moral complacency among the most serious vices (2018, 3113). I think that the seriousness of moral complacency depends on the seriousness of the wrongdoing that is excused as a part of that complacency. The more morally abhorrent the accusations against your friend, the worse it is, morally speaking, to treat them with epistemic partiality. Let's reflect on Carol Reed's *The Third Man* (1949). At first, Martins thinks that Lime is just another petty criminal selling bootlegged cigarettes. In response, he's disappointed in his friend, but he gets over it. And

although it's morally complacent of Martins to excuse Lime's behaviour, I don't think we'd expect him to break off his friendship. When Martins only suspects Lime of bootlegging, the cost of moral complacency is outweighed by the value of their friendship. It's noteworthy that Martins' friendship with Lime is damaged irreparably when he finally grasps the severity of the accusation against Lime. He can't look the other way any longer—he must confront Lime's evil nature head-on. When it becomes salient that the allegations against Lime hold him responsible for the deaths of many children, Martins can be complacent no longer. He stops trying to explain away the evidence presented to him or otherwise exculpate Lime. At this point, he slackens his defensive epistemic standards and comes to accept the allegations against Lime.

Let's consider Stroud's original Sam example (2006), in which he's accused of cruelly breaking someone's heart. When Sam's friends treat Sam with epistemic partiality, it's true that they risk insulting the speaker. But the harm caused by this insult is mild. It shouldn't occupy a very high place in his friends' moral priorities. The thought is that Sam's well-being should occupy a higher place in their moral priorities than the feelings of non-friends. In sections 3.6 and 3.7 of Chapter Two, I argued that it's plausible that our friends are more vulnerable to our harmful beliefs and judgements than non-friends, other things being equal. At least, then, when we have to adjudicate between the testimony of a friend and the testimony of a non-friend on a non-serious matter, it's plausible that our friend's well-being is at greater risk, that is, greater risk of being harmed and at risk of greater harm. Crucially, the harm caused by their treating Sam with epistemic partiality is negligible in comparison with the harm of treating someone accused of racial violence with epistemic partiality. This should occupy a much higher place in Sam's friends' moral priorities. In summary, you shouldn't treat friends with epistemic partiality when they're accused of serious moral misconduct. To do so would be to risk causing serious epistemic injustice in the form of testimonial silencing. Preventing—or at least, not personally causing—epistemic injustice of this kind should be high among your moral priorities.

§5 Summary

In this chapter, I've discussed the a moral objection to treating your friends with epistemic partiality. Treating your friends with epistemic partiality can involve insulting non-friends. In particular, you risk insulting non-friends when you refuse their testimony out of partiality to a friend's well-being. But sometimes, treating your friends with epistemic partiality can perpetuate serious injustice. If a friend is accused of racially or sexually aggravated violence or discrimination, but you refuse the accusation out of loyalty to your friend, then, I argued, it's

plausible that you contribute to the testimonial silencing of the victims of racism or sexism. Our intuitions differ about each of these cases. On the one hand, insulting non-friends is bad, but it doesn't seem right that the risk of insulting a non-friend outweighs the risk of harming a friend. We should treat your friends with epistemic partiality in these circumstances. On the other hand, the risk of perpetuating epistemic injustice seems to outweigh the risk of harming a friend. We shouldn't treat our friends with epistemic partiality in these cases. How do we explain this difference in our intuitions? It's a matter of moral priorities, or so I argued. It's morally permissible that your friends' well-being should matter more to you than the well-being of a non-friend, especially when your friend is at greater risk of harm. But when protecting your friend's well-being requires you to risk perpetuating serious epistemic injustice, it's not morally permissible for you to prioritise your friends' well-being. This suggests that we ought to accept an additional limit on epistemic partiality in friendship, one that's imposed by morality.

Chapter Five

An objection to treating friends with epistemic partiality from the nature of friendship

§1 Introduction

Few people know you better than your best friends, at least according to popular wisdom. And in this instance, popular wisdom seems to have some truth to it. Not only do we share our most intimate secrets with our close friends, but also, there are fewer people better at reading our inner-lives than them. This doesn't appear to be a coincidence either. Part of what bonds us with our close friends is the depth of intimacy we share. We might even think that what our friends know about us should explain why they're our friends. After all, don't you want your friends to love the real you, that is, to love you for the qualities you really possess, rather than, say, some ideal version of you that they've built in their heads?

But if it's so important that our friends know us well, then, arguably, it's unfortunate that we treat our friends with epistemic partiality. What's more, it would be very surprising if we ought to do so, as I argued in Chapter Two. If knowing the facts about our friends and their qualities is so important to our friendships, then how come we should treat our friends with limited epistemic partiality, thereby increasing the chances of missing out on such knowledge? The problem gets worse when you consider that the occasions in which we give our friends the benefit of the doubt are typically situations of personal importance, for instance, when the quality of their character is under question. Here's where the problem arises. To over-simplify things a little: If you value someone's friendship because of their generosity, then it would surely be strange, and perhaps pernicious, if you were to deliberately overlook or underweigh evidence of meanness on their part. One conclusion you might draw from these considerations is that we shouldn't give our friends the benefit of the doubt, since to do so is to risk forming and maintaining friendship with someone based on mistaken conception of them.

In this chapter, I'll consider a final objection to the view that friendship can require epistemic partiality. The objection, owed to Lindsey Crawford (forthcoming), goes like this:

- (1) Your beliefs about your friends should be responsive to the facts about them.

- (2) If your beliefs about your friends should be responsive to the facts about them, then you shouldn't treat your friends with epistemic partiality.
- (3) Therefore, you shouldn't treat your friends with epistemic partiality.

In this thesis I've defended the position that we should treat our friends with limited epistemic partiality. This partiality, or bias, can manifest itself in the way we gather evidence about our friends and in the way we process that evidence. I've also argued that epistemic partiality, so described, can involve the violation of epistemic rationality. The upshot of this is that friendship can require us to hold beliefs against our evidence, and to suspend judgement even though we possess sufficient evidence.

We can think of friendship as a loving relationship (Stump 2011). The thing is, as several philosophers have pointed out, love involves a certain appreciation of the truth about the people we love. Neil Delaney (1996), Simon Keller (2000), Bennett Helm (2010), and Sara Protasi (2016), among others, have each defended views of love according to which to love someone is to value them for their properties—that is, the traits they possess. Even if this view of love isn't correct, it still seems problematic to love someone and all the while be mistaken about what we like about them. We might worry, then, that epistemic partiality is directly inimical to love, so conceived. Suppose this view of love is correct and, accordingly, loving someone involves a certain kind of positive evaluative response to their qualities. If treating someone with epistemic partiality impedes our knowledge of their qualities, either by hiding their flaws or exaggerating their virtues, then in treating someone with epistemic partiality we risk loving them for qualities they don't have, or in spite of flaws that we don't recognise. In other words, epistemic partiality blocks us from loving our friends for what they are, and it blocks our friends from really knowing us too.

I'll argue that epistemic partiality can indeed limit what we know about our friends, but not in ways that diminish the quality of our love for them, and certainly not in ways that block us from being friends with them. I'll question the generality of the second premise, by arguing that oftentimes, when we are actually required to treat our friends with epistemic partiality, we don't come to impede our knowledge of the qualities for which we love them. I'll also note that the character of the epistemic partiality with which we treat our friends is primarily cautious. The main preoccupation is to avoid harming our friends or damaging our union with them. Viewed in this way, the kind of epistemic partiality that friendship can require doesn't lead us into the kind of grave errors of judgement about our friends' characters as the objection supposes. We're not impeded from really knowing our friends.

§2 Your beliefs about your friends should be responsive to the facts about them

I'll start by considering an objection to the claim that friendship requires epistemic partiality made by Crawford (forthcoming), which she calls the 'argument from friendship'. Crawford argues that there's a tension between the nature of friendship and the claim that friendship requires epistemic partiality that can only be reasonably resolved by letting go of the latter claim. First, it's important to draw attention to Crawford's particular conception of epistemic partiality in friendship.

To have doxastically partial beliefs is to have certain generally favorable beliefs and doxastic dispositions that concern one's friends that one would not have concerning relevantly similar non-friends. [...] A good friend's beliefs and doxastic dispositions concerning his friends differ from his beliefs and doxastic dispositions that concern non-friends. And, crucially, they differ in ways that cannot be explained or justified in terms of some difference in the merits that his friends possess, compared with non-friends. (forthcoming, 2)

This is a fairly broad presentation of the epistemological character of epistemic partiality in friendship. I'll briefly highlight a key feature of Crawford's characterisation of the claim that friendship can require limited epistemic partiality. I'll discuss this in the terms of our friends' properties. We come to hold different beliefs and judgements about our friends' properties from those we hold about our non-friends. This difference between our beliefs isn't explained by differences in the properties possessed by our friends and those possessed by our non-friends, though. Rather, what explains the differences are the facts about who's your friend and who isn't. To put it simply, you believe more favourable things about your friends *because they're your friends*.

For Crawford, the partialist, as she calls those who defend the view that friendship can require epistemic partiality (forthcoming, 3), has overlooked an important part of the nature of friendship, that which we discussed in the previous section. Crawford writes:

It is partly constitutive of being a good friend that one's affective and doxastic attitudes toward her friends are responsive to considerations that concern her friend's perceived features, rather than considerations that would appear to make *having* those attitudes sufficiently desirable or choice-worthy to have. (forthcoming, 14)

This is an important step in Crawford's argument. The thought is this: Ideally, you esteem your friends for the properties they hold, or at least, for the properties that you think they actually possess.

According to Crawford, a person's beliefs about her friend should be responsive to the facts about her friend. In the paragraphs that follow, I'll offer broad support for this claim, before I rebut the claim that it entails that friendship cannot require epistemic partiality. To begin, I'll describe what I take to be plausible intuitions about why knowledge matters when it comes to love. Then, I'll explain two ways in which we might shore up these intuitions in philosophical theory. At this point, I'll explain why the notion that knowledge matters in love might make you think that friendship doesn't, and indeed, can't, require even limited epistemic partiality.

Friendship can be understood as a loving relationship. When you love someone, what you know about them (and what you fail to know about them) makes a difference to the quality of that love, or so it seems. Common sense tells us something like the following: You shouldn't be blind to your beloved's failings, nor should you love a false image of their virtues. You should love them for what they are. Consider the last four lines of Shakespeare's Sonnet 130:

I grant I never saw a goddess go;
My mistress when she walks treads on the ground.
And yet, by heaven, I think my love as rare
As any she belied with false compare. (2006, 1074)

Shakespeare's love for his mistress is not impaired by his awareness of her imperfections. Indeed, his love for her seems all the better, all the stronger, all the more *real*, as a result. What's more, it would seem inappropriate to say that he loves her in spite of her flaws. If there's a moral to be drawn from Sonnet 130, it might go something like this: You should love someone for who they really are, warts and all.

The thought at the root of this seems to be that if you really love someone, you shouldn't be ignorant of their flaws. Indeed, a rather stronger but nevertheless plausible reading of this thought is that if you really love someone, you should embrace their flaws as part of them. How should we react when the people we love fail us? How can we love people for who they are, when aspects of their character are bad for them? These questions, of course, can all be applied

to love in friendship too. In any case, I'll not address such considerations in this chapter.⁷⁶ Presumably there are limits to what negative properties you can tolerate in a loved one. This topic has been discussed fruitfully elsewhere.⁷⁷ To briefly illustrate this point, consider what happens in Carol Reed's *The Third Man* (1949) when Martins learns that Lime is accused not of a petty crime, but of racketeering in penicillin, and with fatal consequences. Martins can overlook a little roguishness in Lime's character, but when he learns of Lime's grave moral failings, his love for his friend appears to fail.⁷⁸

So far, I've considered our intuitions about the value of knowing your friends' flaws. But plausibly, knowledge matters when it comes to appreciating someone's positive traits too. In fact, we'll see there are two sides to this. On the one hand, you might hope that the properties for which you love someone are properties that they actually have. On the other hand, from their point of view, they might hope that you love them for properties that they actually have.

Contemporary philosophers are responsive to these intuitions. Stroud writes:

Friendship is in some important sense based on your friend's character and on esteem for his merits. We needn't mean narrowly moral character, or specifically moral merits [...] but I think some such constraint is required in order to respond to the common intuition that we want our friends to *love us for who we are*. (Stroud 2006, 511)

For Keller, the importance of being loved for who goes some way to explain what's good about being loved. We can start to see this if we think about what it's like to be loved. He writes:

If you have a romantic partner, then he is probably forever telling you how wonderful you are. In an ideal relationship, you should have reason to believe him; when other things are going poorly, the fact that he chooses to love you gives you a reason to think that you are not such a bad

⁷⁶ See, for instance, Ward E. Jones (2012) on lover's shame, and Erich Hatala Matthes (2016) on loving someone for who they are in spite of the fact that some aspect of their character is bad for them.

⁷⁷ We might consider Roger Scruton's comment that while he could forgive someone's laziness or selfishness, he would struggle to overlook someone's 'cowardice and viciousness [...] as a murderer' (1986, 240).

⁷⁸ Spoiler: Given the open-ended, cliff-hanger ending of the film, it's not clear whether Martin finally gives up on his friendship with Lime.

person after all. Being loved romantically should give you a reason to feel good about yourself. (2000, 167)

What Keller says about romantic partners here is presumably true about friends too. In an ideal friendship, at least, when your friend pays you a compliment, you hope that their paying you that compliment gives you a reason to feel good about yourself. Kawall (2013, 360) holds a similar view, noting that ‘empty or unwarranted praise brings no satisfaction’. It’s good to receive a compliment from your friend not just because it shows that they care about you, but also because it tells you something good about yourself, or so the thought goes.

Why should we care so much about knowledge in matters of love and friendship? There’s a boring prudential argument which could be made here. If you love someone for qualities they don’t have, or in spite of flaws you’re not aware of, then you might stand to lose out in the long run. Whether or not a prudential argument for the value of knowledge in love and friendship would be successful, it’s not the line of argument that I’ll defend here. I think that, in one way or another, the idea that truth matters in matters of love is central to the way we should understand love itself.

§3 The properties view of love

The straightforward way to defend the idea that we ought to know about our friends’ properties is to adopt a so-called properties view of love. I’ll briefly explain this view and why it might give us cause to doubt that friendship can require epistemic partiality. Now, while the views I consider in this section discuss romantic love, I see no reason why the point can’t be extended to love in general.

Let’s start with Keller’s assessment of the position. He summarises the properties view thus:

The properties view says that the question, “What justifies your choosing to make her the object of your romantic love?” is a sensible question to ask, and that ideal romantic lovers can, in principle, answer it by appealing to a set of the beloved’s properties. (Keller 2000, 164)

According to the properties view, you should love someone because of their properties. This proposal has some initial plausibility. When asked to explain why you love someone, or someone is your friend, it’s natural in your answer to refer to the other person’s properties. Perhaps you love their generosity, their courage, or their intellect. To be more precise, we might say that you

love them *for* their generosity, their courage, or their intellect, since it's the person herself, not their properties, that you love, even though your valuing certain properties explains why you love them, as Delaney (1996) and Keller (2000) have observed.

Not just any properties will do, though. For one thing, it's implausible that loving someone requires maximising your knowledge about them. What's more, some of your beloved's properties are relevant, while others aren't. Luckily, proponents of the properties view have provided a principled way of distinguishing between the relevant and irrelevant properties of the people we love, when it comes to answering the question, 'why do you love her?' The relevant properties are those which relate to the beloved's self-conception, that is, to their identity. Identity here isn't meant in the sense of strict, logical identity, but rather, as what Korsgaard (1996) calls 'practical identity', that is, your understanding of what makes you *you*. As such, a person's practical identity isn't a mysterious thing. Matthes, writing about love, continues,

Some aspects of your practical identity may include character traits, dispositions, or passions; others might include your "social identity," or how you conceive of yourself as belonging to a certain social category. (2016, 243)

There appears to be a degree of consensus among supporters of the properties view on this point. As Delaney writes:

The romantic ideal revolves around the desire to be loved for properties you take to be central to your self-conception; this or something very much like it amounts to the most reasonable interpretation of commonly-used expressions like 'wanting to be loved for who I am'. (1996, 348)

We want to be loved for who we are, or so we say. According to the properties view of love, the best way to understand this is as a declaration of the fact that we want the beliefs and judgements that cause our lover's to love us to be true. We want to be loved for properties that we have prior to being loved.

Several problems enter the picture here, which we might call the problem of substitutability, the problem of trading-up, and the problem of change. I'll briefly describe these problems now by way of introducing Protasi's sophisticated properties view of love (2016). Let's start with the problem of substitutability. If you love someone for their properties, then what happens when you meet their duplicate, that is, someone who holds exactly the same properties as them? Protasi (2016, 220) poses the question: Do you have the same reason to love the *doppelgänger* as you do to love the original? We have a very strong intuition (so it is said) that you

don't. However, it's hard to explain right away how the properties view excludes this possibility. After all, if someone's properties provide your reasons for loving them, then if you encounter someone with the same properties, you are presented with the same reasons for loving the second person as well. The problem of trading-up takes the basic preoccupations that motivate the problem of substitutability and develops them. If we encounter someone who instantiates all the same properties as our beloved, but instantiates them more fully, then surely we have more reasons to love the new acquaintance than we have to love our beloved. But, as Protasi (2016, 220) points out, we have a strong intuition that this conditional is false. The problem of change takes a similar tack (Protasi 2016, 222-223). According to the problem of change, the properties view of love can't account for our intuition that we shouldn't abandon our relationships when their properties change. If love is based on the beloved's properties, how is that our love for them survives changes in their properties?

Protasi (2016) provides a solution to these problems. As Protasi points out in her sophisticated properties view, some of the properties for which you love someone might be relational or historical.⁷⁹ She describes the distinctive senses in which properties can be relational:

As I see it, a property can be relational in two senses. In one sense, a property is relational when its existence or instantiation *depends* on the existence of the relationship [...] We can also talk about relational properties in a wider sense: a property that can be experienced in idiosyncratic and peculiar ways in virtue of an existing relationship. (Protasi 2016, 221)

The properties for which you love someone needn't be simple facts about their character (e.g., 'he is generous, she is courageous, they are intellectual, and so on'). It might not be someone's generosity in itself for which you love them, but for the way they generously helped you out in your time of need. Likewise, you might not love your friend for their courage in itself, but for the courage they showed during some shared tribulation. Yet again, you may not love your friend for their intellect itself, but for the way your intellectual engagement with them has influenced your worldview. If you love them for such relational or historical reasons, then the problem of the *doppelgänger* and the problem of trading-up are deflated somewhat, because the *doppelgänger*, for instance, cannot possess the same properties as your friend. And as for the problem of change: the sorts of historical reasons explained by Protasi cannot change—the friend who generously helped you out in your time of need remains the friend who generously helped you out in your

⁷⁹ Responses of this kind are found in Delaney 1996 and Keller 2000. Amelia Rorty (1986) stresses the importance of the historicity of psychological states in loving relationships.

time of need, regardless of what other changes their character undergoes. If one person's love for their friend is sufficiently based on relational or historical reasons, then, the problem of change doesn't arise.

If the properties view of love is the correct view of love, then it provides us with a fairly straightforward way to explain what explain why you love someone, and why you're justified in loving them, rather than someone else. Simply, it's because of your response to the properties they have. As we've seen, these properties could be relational or historical as well as internal or ahistorical: that helps us to overcome the problems of substitutability or trading-up.

While the properties view of love appears to provide us with a straightforward way to understand the importance of knowledge in our close, personal relationships. In mitigation, the view faces several objections.⁸⁰ But even if the properties view of love is incorrect, it's still plausible that to some extent it should trouble us if we esteem our friends for properties they don't have, or in spite of the fact they possess grave character flaws that we've overlooked. Now I'll explain why these considerations appear to give us a reason to doubt that friendship can require epistemic partiality.

§4 If your beliefs about your friends should be responsive to the facts about them, then you shouldn't treat your friends with epistemic partiality

The central thought to the previous section was that what you believe about them should be responsive to the facts about them. Focussing on esteem, we considered the suggestion that you should esteem your friends for properties that they actually possess. Now let's consider an alternative, namely, that you esteem your friends for properties that you judge it would be fruitful for you to believe that they hold. To esteem your friend for such reasons would seem inappropriate. Surely, so the thought goes, if you esteem your friend for her generosity, it should be because she's actually generous, and not merely because you judge that it will benefit your friendship to believe that she's generous. Why is this? Well, according to Crawford, one reason it seems inappropriate to respond to such considerations is that your esteem for your friend wouldn't be appropriately responsive to the friend's actual qualities. She writes:

In order for the good friend's attitudes of esteem for her friend to be appropriately responsive to the qualities that the friend has, certain things must be true about the good friend's beliefs. The

⁸⁰ There are notable objections from Velleman (1999), Kolodny (2003) and Frankfurt (2004).

good friend must believe that the friend has the various qualities that merit the esteem she has for her friend. You esteem your friend's kindness, for example, when you think she *really is* kind. And your belief about your friend's kindness is itself appropriately responsive to your friend when it is based on your judgment that the friend in fact has those qualities. You do not, by contrast, believe that the friend is kind *because* believing that she is will make it easier for you to esteem her. (2017, 13)

The notion of appropriate responsiveness enters the stage at this point. Crawford explains the notion of appropriate responsiveness in terms of the difference between state- and object-given reasons for belief (2017, 13). I'll quickly describe this distinction before continuing with my discussion of Crawford's argument.

You have object-given reasons to believe that you'll successfully jump across the ravine if, for instance, you have evidence according to which it's more likely that you'll successfully jump across the ravine than not. Such evidence might include the facts about your track record and your assessment of the difficulty of the jump. Where object-given reasons for believing that p relate to the truth of the proposition that p , state-given reasons have to do with the practical rationality of believing that p . You have state-given reasons to believe that you'll jump across the ravine, if, for instance, you believe that the more confident you are that you'll successfully jump across the ravine, the more likely you are to successfully jump across the ravine (assuming, of course, that you actually want to.)

For the remainder of this chapter, I'm going to assume that evidentialism is true. Crawford doesn't make this assumption. By discussing epistemic partiality in terms of object- and state-given reasons, she leaves open the question of which theory of epistemic justification is correct. If evidentialism is correct, only evidence that p provides object-given reasons to believe that p . On the assumption that evidentialism is correct, then, we can interpret Crawford's claim this: It's constitutive of friendship that what you believe about your friends, you should believe on the basis of your evidence. But according to Crawford's account of epistemic partiality, the following conditional is true: If we should treat our friends with limited epistemic partiality, then friendship can require us to believe not on the basis of our evidence. (Rather, it requires us to believe for state-given, or, as I've characterised epistemic partiality, moral, reasons.) By *modus tollens*, we can conclude that it's not the case that we should treat our friends with epistemic partiality.

Let's consider some of the cases that we've already discussed in this thesis. Take Keller's Poetry Reading example (2004) for starters. Let's start by thinking about the reasons why Eric likes Rebecca. Well, according to the argument under consideration, if Eric esteems Rebecca for

her poetic skill, it should be because she really has that skill. So what kinds of considerations provide the appropriate kinds of reasons for him to esteem Rebecca for her poetic skill? Only his evidence can provide such reasons. He gains this evidence by going to her poetry reading and paying attention to her poetry. If the evidence he ends up with at the end of the reading supports the judgement that Rebecca's a skilled poet, then it's reasonable for him to esteem her for her poetic skill. If his evidence doesn't support such a judgement, then it's not. As such, state-given reasons, such as reasons about their friendship status, or pragmatic reasons more generally, such as reasons about the harm that his failure to judge her poetry positively, are excluded from consideration.

Now let's consider what Rebecca might expect from Eric's judgements about her. In this case, let's suppose, Rebecca wants to be esteemed by her friends for the skill she displays in her performance. Why does she want to be esteemed as such? The answer, according to the argument presented here, is simply that she wants to be a skilled poet. Perhaps, this is an important part of her practical identity, or an important ground-project of hers. Given the assumption that what matters to Rebecca is that she actually is a skilled poet, we can explain for what reasons she would expect Eric to believe that she's a skilled poet. Simply, she would want him to believe that she's a skilled poet because his evidence shows that she is one. Indeed, as Kawall (2013) suggests, friends are in a good source of this kind of criticism because we can feel assured that they have our best interests at heart.

Whereas the Poetry Reading example focusses on how we form positive beliefs about our friends, Stroud's Sam example (2006) focusses on how we withhold from forming harmful, often negative, beliefs about them. Reflection on this case also seems to offer support for Crawford's objection. Suppose that Sam's friends' esteem for him would be damaged if they came to believe that he was cruelly dismissive of his sexual partners. As it happens, they believe that he's not cruelly dismissive of his sexual partners. To put it simply, if they're good friends, their judgements about Sam's character are appropriately responsive to his actual character. The sharp end of this point is that if Sam is in fact cruel, they shouldn't be friends with him. The best and indeed the only way for Sam's friends to avoid the possibility that they esteem him for his lack of cruelty when he is, in fact, cruel, is for them to believe according to their evidence. As such, state-given reasons, such as reasons about their friendship status, or pragmatic reasons more generally, such as reasons about the harm that his failure to judge her poetry positively, are excluded from consideration.

For this reason, according to Crawford, it's not the case that you should treat your friends with limited, epistemic partiality. In the section that follows, I'll respond to Crawford's objection.

§5 A response to the objection from the nature of friendship

First, it's important to draw attention to the difference between the Crawford's conception of epistemic partiality in friendship and my own. Crawford's presentation of the character of epistemic partiality in friendship is rather stronger than the view I've defended in this thesis. Where she argues that epistemic partiality in friendship requires believing for state-given, pragmatic reasons, I've argued that epistemic partiality may also include behaviours such as: (1) gathering evidence in biased ways; (2) processing your evidence in biased ways. None of these forms of epistemic partiality necessarily involve believing for state-given reasons. State-given, pragmatic reasons appear to play a different role.

Consider the claim that friendship can require us to gather evidence about our friends in biased ways. It seems as if our reasons for gathering evidence in biased ways could be state-given reasons. In *The Third Man* example, plausibly, Martins responds to state-given reasons, the fact that he's friends with Lime, when he searches for exculpatory evidence. But that doesn't mean he believes for state-given, pragmatic reasons. The beliefs that Martins comes to hold are informed by his evidence, and his evidence alone. So it's not the case that the epistemic partiality with which he treats Lime involves belief for state-given, pragmatic reasons.

Now consider the claim that epistemic partiality in friendship requires processing your evidence in biased ways. One of the ways in which we treat our friends with epistemic partiality is to require more evidence before we believe harmful things about them. We can understand this as a matter of withholding belief. But, as Schroeder points out (2012, 277), there's a significant difference between withholding belief for non-evidential, pragmatic reasons and forming beliefs for such reasons. He writes that '[t]o withhold is to not make up your mind, to have formed no belief. Consequently, any disadvantage of forming beliefs—of making up your mind—is potentially a reason to withhold' (2012, 277). But evidence can't determine whether it's rational for you to withhold belief, since, as Schroeder writes, evidence is relevant only for determining whether it's the case that p , and not for deciding whether to make up your mind about whether p . So, when determining whether to withhold judgement with respect to the proposition that, say, your friend Sam is a scoundrel, the facts about whether he's in fact a scoundrel aren't relevant, but the facts about the consequences of holding that belief are. If you

end up believing (or disbelieving) that Sam is a scoundrel, then you believe in response to your evidence. But you can't withhold judgement for evidential reasons. Rather, when you withhold judgement, you do so for pragmatic reasons, but in a rationally acceptable way.

For these reasons, I think that Crawford's characterisation of epistemic partiality in is incomplete. Nevertheless, her objection has other merits, and it would be churlish not to take them into consideration. In the paragraphs that follow, I'll discuss Crawford's view of the nature of friendship.

The main force of Crawford's objection comes from her claim that it's constitutive of friendship that the beliefs and judgements we form about our friends' properties are responsive to their actual properties. The thought is that esteem, to be real esteem, genuine esteem, must be appropriately responsive to qualities that actually merit esteem. Kawall's complaint (2013), namely, that false compliments offer little comfort, seems sensitive to the same concerns as Crawford's: If what your friends believe about you is epistemically unjustified, then you don't benefit meaningfully from their believing it. Now I'll assess this claim.

I do not agree with Kawall and Crawford on this point. It's not clear that we want the beliefs of our friends and close, personal relations to be based on object-given, but not state-given, reasons. In fact, it's plausible that oftentimes we demand that people believe things about us for state-given reasons. Let's reflect for a moment on some fairly common utterances:

'It's nice that she thinks so highly of me, even if it's only because she's my friend.'

'I can't believe he took those rumours seriously; we're supposed to be friends!'

'Please believe me; we're friends.'

'If you're my friend, you won't listen to him.'

In balance, these types of utterance might just be idiosyncratic ways of signposting object-given, that is to say, evidential reasons for belief. Following this line of reasoning, the reason someone reminds you that you're friends when they want you to believe them is because they take the fact that you're friends to be evidence that you can be trusted to tell them the truth. The implication is that, given that we're friends, it's improbable that I'm being insincere in my utterance.

But I think an equally plausible way of reading these utterances is as appeals to believe for state-given reasons. At least, in each example, the speaker seems to assent to the notion that their friend believe for state-given reasons rather than for object-given ones. What does this show us? Well, these early reflections give us reasons to doubt that what Crawford describes as a condition of friendship actually prescribes that our beliefs about our friends' properties are

appropriately responsive to her actual properties. Keller describes his expectation that his running friends form beliefs about him in just this way.

Among the things I want from my running friends is feedback that will keep me motivated. Sometimes that purpose is served by a completely accurate assessment of my performance and my prospects. But when I am fighting against age and injury and poor form, sometimes I just need to hear something positive, and it is valuable to have friends who see the best in me and believe the best about my potential to improve. I know that as my friends, their opinions are biased and I should interpret their feedback in that context. And I know that if I were a stronger, less needy person, then I would have no interest in being encouraged and supported; I could take care of that myself. But still, being the person I am, those are some of the good things I get from having friends. (2018, 27)

This is just one example, but I think it shows us something interesting about what we expect from our friends. It shows us how, when we have expectations about our friends' beliefs, we're not always concerned about whether their beliefs are true or justified. It doesn't always matter to us whether our friends' beliefs are candidates for knowledge, so to speak. Rather, sometimes we want our friends to hold certain beliefs because of the consequences of their holding that belief—for instance, the encouragement it gives us to know that they see the best in us.⁸¹ Friendship is just more complicated than this. To echo Basu (2018, 4) it's incorrect to say that the our we always want the truth from our friends. No, friendships are much more complicated than that. Sometimes we want our friends to back us fully, which in this context means to have unwavering confidence in our good character and competence. Some other times, perhaps before we put ourselves in publicly embarrassing situations, we'd prefer that a friend accurately perceived our shortcomings.

Let's pause to consider a possible response. You might object that all this shows is that friendship can require limited epistemic partiality, but not with respect to the reasons for which we esteem our friends. The amended view would then be that you should treat friends with epistemic partiality, but only with respect to beliefs that don't explain why we esteem each other,

⁸¹ There are also cases in which we consider belief not for object-given reasons eminently virtuous. Ryan Preston-Roedder (2013) has written instructively on the concept of faith in humanity. Someone who has faith in humanity 'tends to believe in people, trust in them, make presumptions in their favor, or see them in a favorable light, morally speaking' (Preston-Roedder 2013, 666).

that is to say, why we like each other. This would seem to eviscerate the thesis that you should treat friends with epistemic partiality, since it would limit that partiality to a narrow group of beliefs about your friends. In particular, the narrow group of beliefs would be those which (a) aren't the reasons why you esteem your friend and/or feel a bond with them, but (b) would harm your friend, if you came to hold them.

But this is to concede too much. The thought behind the objection was that we shouldn't like our friends for qualities they don't have. But of course, even the someone who claims that friendship can require epistemic partiality would accept this claim! We shouldn't think of epistemic partiality in friendship as the practice of painting a false picture of our friends for the purpose of building a self-perpetuating esteem-loop. Rather, treating our friends with epistemic partiality is a matter of preparing for a particular risk, namely, the risk of incorrectly believing (or disbelieving) something about your friend which harms them.

In fact it shouldn't surprise us to think of friendship as involving incomplete or partial knowledge. C.S. Lewis describes how our knowledge of our friends is informed by their interactions with our other friends:

[I]f, of three friends (A, B, and C), A should die, then B loses not only A but 'A's part in C', while C loses not only A but 'A's part in B'. In each of my friends there is something that only some other friend can fully bring out. By myself I am not large enough to call the whole man into activity; I want other lights than my own to show all his facets. Now that Charles is dead, I shall never again see Ronald's reaction to a specifically Caroline joke. Far from having more of Ronald, having him 'to myself' now that Charles is away, I have less of Ronald. Hence true Friendship is the least jealous of loves. (1960, 73-74)

Lewis isn't writing specifically about our knowledge of our friends here, but there's still an important moral to be drawn from this passage. The thought seems to be this: It's ordinary not to have complete knowledge of your friends' characters. In fact, sometimes we can only access certain parts of our friends' characters in the company of some third friend, but even then, we only get glimpses: only a third friend can 'fully bring out' some qualities in your friend. The point, which Lewis expresses here, is that your having incomplete knowledge of a friend is not a problem for your friendship with them. It's not the case that A's friendship with B is defective to the extent that there are aspects of A's personality that B can only properly encounter when C is around. Suppose A values artistic excellence, and only reveals this trait around C. Now, B may learn that A values artistic excellence from the time they spend together with C, without ever learning much about how good A is at art. This hardly diminishes the quality of their friendship.

However, if B were mistaken about whether A values art, their friendship would seem diminished. A lack of knowledge about his friend's values is more damaging than a lack of knowledge about his friend's properties. This moral isn't supposed to generalise. We can accept that there are important things to learn about our friends' properties too. Rather, the point is that knowledge of our friend's properties isn't the only thing that matters for the flourishing of our friendships with them.

§6 Summary

I have argued that it's not the case that it's constitutive of friendship that your beliefs about your friends should be responsive to the facts about them, at least, not in a way that means that we must only form beliefs and judgements about them for object-given, evidential reasons. Even so, I continued to argue that even if your beliefs about your friends should be responsive to the facts about them, then friendship doesn't require epistemic partiality. That's because the best explanation of why your beliefs about your friends should be responsive to the facts about them limits the reach of the objection to those beliefs about which you esteem someone. Finally, I argued that it's not even clear that rationally permissible esteem for a friend completely excludes epistemic partiality, because it's neither necessary nor always possible to have complete knowledge of someone's properties to esteem them.

Conclusion

§1 The epistemology of epistemic partiality in friendship

We should treat our friends with limited epistemic partiality, or so I argued in this thesis. We should give them the benefit of the doubt when they face criticism from others, and we should look on the bright side when considering their virtues. These practices—giving the benefit of the doubt and looking on the bright side—are two of many varieties of epistemic partiality with which we often treat our friends.

This thesis was divided into two parts. In the first part I presented the positive case for limited epistemic partiality in friendship. In the second part of the thesis I attempted to rebut three objections to the position defended in the first part. In the first chapter, I introduced the phenomenon of epistemic partiality. Inspired by Sarah Stroud's descriptivist methodology (2006), I reflected on the experience of treating others with epistemic partiality. I argued that epistemic partiality is commonplace in close, personal relationships: We often give our friends the benefit of the doubt or look on the bright side when we consider their character and conduct. What's more, I claimed, we sometimes feel as if we should treat our friends and close, personal relations in this way. That was one of the lessons of Shakespeare's *Othello*: It seems as if Othello wrongs Desdemona when he fails to give her the benefit of the doubt in the face of accusations of infidelity.

One of the goals of this thesis was to offer an analysis of epistemic partiality in friendship. It's true that we comfortably speak of giving the benefit of the doubt and looking on the bright side. But what exactly is it to give someone the benefit of the doubt or to look on the bright side? How common are these behaviours? In Chapters One and Two, I attempted to provide a deeper analysis of the behaviours that constitute epistemic partiality in friendship. In Chapter One, I identified two varieties of epistemic partiality in friendship. Epistemic partiality, I suggested, involves two distinctive behaviours. First, we show epistemic partiality in the way we gather evidence about our friends. One way in which we treat a friend with epistemic partiality at this level is that we gather evidence which supports a more favourable conclusion about their character and conduct. Another way we gather evidence in epistemically partial ways is by declining to gather evidence which would confirm a less favourable conclusion about them. Second, we show epistemic partiality in the way we process our evidence about our friends. For example, we withhold judgement about our friends when we possess evidence that would support a less favourable interpretation of their character or conduct. Alternatively, we demand

less evidence before we're able to form favourable beliefs about a friend than we would demand to believe that same thing about a stranger.

In Chapter Two, I argued that we should treat our friends with epistemic partiality, albeit within principled limits. The starting point of my argument is the thought that what you believe about your friends—and what you don't believe about them—can harm them. In defence of this premise, I explored the burgeoning epistemology of harmful beliefs. There's no single way in which beliefs can harm. Rather, several, compatible epistemologies of harmful beliefs provide a richer way to understand the harms that our beliefs can cause. Our beliefs can conflict with our friends' desires and cause them pain (Keller 2018), thereby harming them. We harm our friends when our judgements tarnish them with undue criticism, even if that criticism is justified from the believer's point of view. The way we form beliefs about our friends can treat amount to treating them as objects rather than as persons (Basu 2019), and this too constitutes a wrong against them. Finally, we can fail to show adequate goodwill in the way that we form beliefs about them (Marušić and White 2018). What's more, I argued, our friends are especially vulnerable to harm from our beliefs for two reasons, namely, the importance of our beliefs in the formation of their self-esteem and the fact that friendship involves a commitment to care which, when broken, compounds the harms of harmful beliefs. That's why we don't have to treat everyone with epistemic partiality.

What's this got to do with friendship? Well, as I claimed, friendship is a caring relationship, according to almost all philosophical accounts of friendship. If you count someone as your friend, then you're committed to caring about them. At a minimum, caring about someone involves holding certain desires about their well-being. Granted that beliefs can harm and the fact that friendship involves a special commitment of care towards your friends, we should be especially concerned by the risk that our beliefs harm our friends. For instance, if, as in Stroud's Sam example (2006), Sam's friends accepted the accusation against Sam, then it's plausible that they, that is, their beliefs, would harm him. The judgements they would come to hold constitute a kind of damage to his reputation among his friends—damage that can be mitigated only if they treat him with epistemic partiality. Likewise, in the Keller's Poetry Reading example, if Eric didn't look on the bright side of Rebecca's performance, then he seems to harm her too. In reviewing her performance from a cool and impersonal standpoint, he fails to show her the level of goodwill that's appropriate given their relationship. After considering these and several other putative examples of epistemic partiality in friendship, I concluded that oftentimes the only way to comply with your commitment to care for our friends' well-being is to treat them with epistemic partiality. If we desire a friend's well-being and epistemic partiality is the only way

to get it, and we don't possess stronger pragmatic reasons against treating that friend with epistemic partiality, then we should treat the friend with epistemic partiality.

I acknowledged that there must be limits to the extent to which treating our friends with epistemic partiality can actually benefit them. While treating our friends with epistemic partiality can protect and promote their well-being, it can harm them too. Sometimes it's important for us to know the truth about our friends' character and conduct. This shows us that there must be limits to the extent to which we should treat our friends with epistemic partiality.

In the second part of this thesis, I responded to three objections to the thesis that we should treat our friends with epistemic partiality. First, in Chapter Three, I responded to what I called the epistemological objection to the argument from care. The epistemological argument went like this: (1) Epistemic partiality in friendship is epistemically irrational. (2) If epistemic partiality in friendship is epistemically irrational, then it's irrational, all things considered. (3) Therefore, epistemic partiality in friendship is irrational, all things considered. First, I explained why the objector might think that epistemic partiality is epistemically irrational, at each of the levels of evidence gathering and evidence processing. Then, I explained why they might think that that would imply that epistemic partiality is irrational, all things considered. Then, I explained why it's not so obvious that epistemic partiality is epistemically irrational. Drawing on Kawall (2013) and Goldberg (forthcoming), I suggested some reasons for thinking that apparently partial ways of gathering evidence are, in fact, epistemically rational. Then, I suggests some reasons for thinking that epistemic partiality in evidence processing isn't epistemically irrational either. First, I considered Hawley's (2014) proposal that, if permissivism is true, then the first premise of the argument is unmotivated. Then, I considered Schroeder's (2018a) argument that, assuming pragmatic intellectualism is true, treating people with epistemic partiality by suspending judgement in spite of your evidence is not epistemically irrational. For argument's sake, I remained agnostic about whether epistemic partiality is, in fact, epistemically irrational. Even if epistemic partiality in friendship is epistemically irrational, I argued, that doesn't imply that it's all things considered irrational. For one thing, it's not clear that epistemic reasons can override moral reasons, since moral and epistemic reasons seem to be incommensurable. But even if epistemic reasons can override moral reasons, it's not clear that the moderate epistemic irrationality involved in treating your friends with epistemic partiality is so costly that it overrides what moral reasons we possess for treating with limited epistemic partiality. For these reasons, the epistemological objection fails.

Then, in Chapter Four, I considered a moral objection to the thesis. In particular, the objection maintains that it's immoral to treat our friends with epistemic partiality. In setting up

the objection, I focussed on cases of epistemic partiality in which someone refuses a speaker's testimony in order to avoid forming a belief that would harm their friends. I have suggested that refusing someone's testimony can constitute an insult towards them, at least when you don't possess good epistemic reasons for refusing it. Insulting people is regrettable, but that's not the full extent of the harm that epistemic partiality can cause. In certain circumstances, I argued, epistemic partiality in friendship might also perpetuate testimonial injustice. That's because refusing someone's testimony when they report, for instance, racially- or sexually-aggravated violence or discrimination, and doing so for reasons of friendship, is likely to satisfy the conditions of testimonial silencing. This can have the knock-on effect of smothering further potential speakers. If someone knows that you'll not accept criticism of your friends, then it's likely that they won't even attempt to tell you when they believe your friend is guilty of serious moral misconduct. In the end, what this shows us is that there are additional, moral limits to the extent to which we should treat our friends with epistemic partiality. We shouldn't treat your friend with epistemic partiality when they're accused of serious misconduct. To do so, I argued, would be to act with an objectionable degree of moral complacency.

Finally, in Chapter Five, I considered an objection levelled by Lindsay Crawford, according to whom epistemic partiality is incompatible with friendship. According to Crawford (forthcoming), it's constitutive of friendship that our beliefs about our friends are based on object-given reasons. Exploring the properties view of love, I considered some additional reasons why we might think that friendship requires us to be responsive. A friendship is in some sense defective to the extent that what one friend believes about the other is based on state- rather than object-given reasons, at least according to objection. I think that this objection fails. Friendship doesn't require belief for object-given reasons—in fact, oftentimes we would want our friends to believe us for state-given reasons. When it's important that you take my word for something, I want you to believe me because you're my friend, not because you don't happen to possess overwhelming, non-testimonial reasons to accept my testimony. What's more, I claimed, we sometimes think that friends are commendable for their unwavering loyalty, even when such loyalty involves epistemically irresponsible behaviour.

And that brings us to the conclusion. We should treat our friends with epistemic partiality because we care about them. Such epistemic partiality may be epistemically irrational—I remained agnostic on that controversial point. However, our care itself imposes limits on the extent to which we should treat them with epistemic partiality. Epistemic partiality in friendship is morally risky—if we're unwavering in our loyalty to our friends then we may risk causing greater injustices. It's reasonable to conclude that there are both practical limits and moral limits

to the extent to which we should treat our friends with epistemic partiality. Nevertheless, I conclude that when your friend's well-being is at stake, you treat them with epistemic partiality.

§2 Further research

So what next, for the epistemology of epistemic partiality? I've already argued that in situations where our beliefs harm, we must consider whether we should treat the victim of that harm with epistemic partiality. If we fail to accept the reports of victims of sexual or racial discrimination or violence, we risk causing them serious harm—compounding the damage that has already been done. Should we treat the self-reporting victims of sexual or racial discrimination or violence with epistemic partiality when they report the crimes that have been committed against them?⁸²

It's plausible that we harm the victim of racial or sexual discrimination or violence if we refuse their testimony, when they report the crimes committed against them. The victim of racial or sexual discrimination or violence often suffer from a kind of shame, where they (incorrectly) judge that they deserve rejection. If we reject their testimony when they report what happened to them, we compound this sense of shame by rejecting them all over again. To avoid causing additional harm to the victim, then, we must accept their testimony. The problem is that sometimes epistemic propriety demands that we reject the victims testimony. For example, when the speaker's credibility is unknown to us, and what they allege is improbable, given what other evidence we have, it may be epistemically unreasonable to accept their testimony. Epistemically speaking, we should not accept their testimony. If that's the case, then acting in accordance with epistemic propriety would cause us to harm the victim by refusing their testimony. This is where epistemic partiality becomes important. I think it's plausible that we should treat the victim here with epistemic partiality by accepting their testimony in spite of the fact that it's not epistemically rational to do so. We do so to protect the victim's well-being.

Now, several objections will spring to mind. First, one might doubt that accepting someone's testimony is necessary or sufficient for protecting them from harm. Second, you might doubt that it's epistemically unreasonable to accept the victim's testimony in such cases. Third, you might be concerned about other responsibilities we have, for instance, the responsibility not to believe false accusations. Presumably, institutional responsibilities make a difference here too. What capacity you're in when you receive the report could conceivably make a difference to whether you should treat the victim with epistemic partiality. It's perhaps less

⁸² This is a line of research that I'm currently developing with David Eford.

appealing to our intuitions that the jurors in a criminal case ought to treat a victim with epistemic partiality. These are all issues that will have to be resolved elsewhere. Nevertheless, I think that this is a promising line of future research which shows how the epistemology of epistemic partiality could develop into a morally important area of philosophical research.

References

- Annis, D. B. (1987). The meaning, value, and duties of friendship. *American Philosophical Quarterly* 24(4), 349-356.
- Alston, W. P. (1985). Concepts of epistemic justification. *The Monist* 68(2), 57-89.
- Alston, W. P. (1988). The deontological conception of epistemic justification. *Philosophical Perspectives* 2, 257-299.
- Aristotle (2009). *The nicomachean ethics*. Translated by David Ross. New York, NY: Oxford University Press.
- Arpaly, N. and Brinkerhoff, A. (2018). Why epistemic partiality is overrated. *Philosophical Topics* 46(1), 37-51.
- Badhwar, N. K. (1987). Friends as ends in themselves. *Philosophy and Phenomenological Research* 48(1), 1-23.
- Badhwar, N. K. (ed.) (1993). *Friendship: a philosophical reader*. Ithaca, NY: Cornell University Press.
- Baehr, J. (2011). *The inquiring mind: on intellectual virtues and virtue epistemology*. Oxford: Oxford University Press.
- Basu, R. (2018). Can beliefs wrong? *Philosophical Topics* 46(1), 1-17
- Basu, R. (2019). What we epistemically owe to each other. *Philosophical Studies* 176(4), 915-931.
- Basu, R. (forthcoming). The wrongs of racist beliefs. *Philosophical Studies*, 1-19.
- Basu, R. and Schroeder, M. (2019). Doxastic wronging. In B. Kim & M. McGrath (eds.), *Pragmatic encroachment in epistemology*. London: Routledge, 181-205.
- Bentham, Jeremy (1789/1907). *An introduction to the principles of morals and legislation*. Oxford: Clarendon Press.
- BonJour, L. (1985). *The structure of empirical knowledge*. Cambridge MA: Harvard University Press.
- Burge, T. (1993). Content preservation. *Philosophical Review* 102(4), 457-488.
- Bradley, B. (2009). *Well-being and death*. Oxford: Oxford University Press.
- Clifford, W. K. (1877). The ethics of belief. *Contemporary Review* 29, 289-309.
- Cocking, D. and Kennett, J. (2000). Friendship and moral danger. *Journal of Philosophy* 97(5), 278-296.
- Cohen, G. A. (2000) *If you're an egalitarian, how come you're so rich?* Cambridge MA: Harvard University Press.
- Cooper, J. M. (1977). Friendship and the good in Aristotle. *Philosophical Review* 86(3), 290-315.
- Crawford, L. (forthcoming). Believing the best: on doxastic partiality in friendship. *Synthese*, 1-19.

- Crenshaw, K. (1991). Mapping the margins: intersectionality, identity politics, and violence against women of color. *Stanford Law Review* 43, 1241-1299.
- Crisp, R. (2006). Hedonism reconsidered. *Philosophy and Phenomenological Research* 73(3), 619-645.
- Cross, A. (manuscript). The puzzle of partiality.
- David, M. (2001). Truth as the epistemic goal. In Matthias Steup (ed.), *Knowledge, truth, and duty*. New York: Oxford University Press, 151-169.
- Davis, E. (2016). Typecasts, tokens, and spokespersons: a case for credibility excess as testimonial injustice. *Hypatia* 31(3), 485-501.
- Delaney, N. (1996). Romantic love and loving commitment: articulating a modern ideal. *American Philosophical Quarterly* (33)4, 339-356
- Dogramaci, S. and Horowitz, S. (2016). An argument for uniqueness about evidential support. *Philosophical Issues* 26(1), 130-147.
- Dotson, K. (2011). Tracking epistemic violence, tracking practices of silencing. *Hypatia* 26(2), 236-257.
- Dougherty, T. (ed.) (2011). *Evidentialism and its discontents*. Oxford: Oxford University Press.
- Fantl, J. and McGrath, M. (2002). Evidence, pragmatics, and justification. *Philosophical Review* 111(1), 67-94.
- Fantl, J. & McGrath, M. (2009). *Knowledge in an uncertain world*. New York: Oxford University Press.
- Faulkner, P. (2018). Giving the benefit of the doubt. *International Journal of Philosophical Studies* 26(2), 139-155.
- Feldman, F. (2004). *Pleasure and the good life: concerning the nature, varieties, and plausibility of hedonism*. Oxford: Clarendon Press.
- Feldman, R. & Conee, E. (1985). Evidentialism. *Philosophical Studies* 48(1), 15-34.
- Feldman, R. (2000). The ethics of belief. *Philosophy and Phenomenological Research* 60(3), 667-695
- Feldman, R. (2007). Reasonable religious disagreements. In L. Antony (ed.), *Philosophers without gods: meditations on atheism and the secular life*. New York, NY: Oxford University Press, pp. 194-214.
- Fletcher, G. (2015). Objective list theories. In G. Fletcher (ed.), *The routledge handbook of philosophy of well-being*. London: Routledge, pp. 148-160.
- Finnis, J. (1980). *Natural law and natural rights*. Oxford: Clarendon Press.
- Frankfurt, H. (1982). The importance of what we care about. *Synthese* 53(2), 257-272
- Frankfurt, H. (1998). Duty and love. *Philosophical Explorations* 1(1), 4-9.

- Frankfurt, H. (1999). Autonomy, necessity, and love. In *Necessity, Volition, and Love*, Cambridge: Cambridge University Press, pp. 129-41.
- Frankfurt, H. (2004). *The reasons of love*. Princeton, NJ: Princeton University Press.
- Fricker, E. (1994). Against gullibility. In A. Chakrabarti and B. K. Matilal (eds.), *Knowing from words*. Kluwer Academic Publishers.
- Fricker, M. (1998). Rational authority and social power: towards a truly social epistemology. *Proceedings of the Aristotelian Society* 98(2), 159-177.
- Fricker, M. (2007). *Epistemic injustice: power and the ethics of knowing*. Oxford: Oxford University Press.
- Gendler, T. S. (2011). On the epistemic costs of implicit bias. *Philosophical Studies* 156(1), 33-63.
- Grice, H. P. (1989). *Studies in the way of words*. Cambridge MA: Harvard University Press.
- Grunebaum, J. O. (2005). Fair-weather friendships. *Journal of Value Inquiry* 39(2), 203-214.
- Goldberg, S. C. (forthcoming). Against epistemic partiality in friendship: value-reflecting reasons. *Philosophical Studies*, 1-22.
- Hall, R. J. & Johnson, C. R. (1998). The epistemic duty to seek more evidence. *American Philosophical Quarterly* 35(2), 129-139.
- Hawley, K. (2014). Partiality and prejudice in trusting. *Synthese* 191(9), 2029-2045.
- Hawthorne, J. (2003). *Knowledge and lotteries*. Oxford: Oxford University Press.
- Hazlett, A. (2017). On the special insult of refusing testimony. *Philosophical Explorations* 20(1), 37-51.
- Heathwood, C. (2006). Desire satisfactionism and hedonism. *Philosophical Studies* 128(3), 539-563.
- Heathwood, C. (2016). Desire-fulfilment theory. In G. Fletcher (ed.), *The routledge handbook of the philosophy of well-being*. Routledge, pp. 135-147.
- Helm, B. (2010). *Love, friendship, and the self: intimacy, identification, and the social nature of persons*. Oxford University Press.
- Hewitt, S. (2010). What do our intuitions about the experience machine really tell us about hedonism? *Philosophical Studies* 151(3), 331-349.
- Hieronymi, P. (2005). The wrong kind of reason. *Journal of Philosophy* 102(9), 437-457.
- Hieronymi, P. (2001). Articulating an uncompromising forgiveness. *Philosophy and Phenomenological Research* 62, 529-555
- Holton, R. (1994). Deciding to trust, coming to believe. *Australasian Journal of Philosophy* 72(1), 63-76
- Hookway, C. (2008). Questions, epistemology, and inquiries. *Grazer Philosophische Studien* 77, 1-21.
- Horowitz, S. (2014). Immoderately rational. *Philosophical Studies* 167(1), 41-56.

- Isserow, J. (2018). On having bad persons as friends. *Philosophical Studies* 175(12), 3099-3116.
- James, W. (1987/1992). The will to believe. In G. Myers (ed.), *William James: writings 1878-1899*. USA: The Library of America.
- Jones, W. E. (2012). A lover's shame. *Ethical Theory and Moral Practice* 15(5), 615-630.
- Kagan, S. (1992). The limits of well-being. *Social Philosophy and Policy* 9(2), 169-189.
- Kawall, J. (2013). Friendship and epistemic norms. *Philosophical Studies* 165(2), 349-370.
- Keller, S. (2000). How do I love thee? Let me count the properties. *American Philosophical Quarterly* 37(2), 163-173.
- Keller, S. (2004). Friendship and belief. *Philosophical Papers* 33(3), 329-351.
- Keller, S. (2013). *Partiality*. Princeton, NJ: Princeton University Press.
- Keller, S. (2018). Belief for someone else's sake. *Philosophical Topics* 46(1), 19-35.
- Kelly, T. (2003). Epistemic rationality as instrumental rationality: a critique. *Philosophy and Phenomenological Research* 66(3), 612-240.
- Kleinig, J. (2014). *On loyalty and loyalties: the contours of a problematic virtue*. New York, NY: Oxford University Press.
- Kolodny, N. (2003). Love as valuing a relationship. *Philosophical Review* 112(2), 135-189.
- Kolodny, N. (2010). Which relationships justify partiality? General Considerations and Problem Cases. In Brian Feltham and John Cottingham (eds.), *Partiality and impartiality: morality, special relationships, and the wider world*. Oxford: Oxford University Press.
- Koltonski, D. (2016). A good friend will help you move a body: friendship and the problem of moral disagreement. *Philosophical Review* 125(4), 473-507.
- Korsgaard, C. M. (1996). *The sources of normativity*. Cambridge: Cambridge University Press.
- Kvanvig, J. (2005). Truth is not the primary epistemic goal. In M. Steup and E. Sosa (eds.), *Contemporary debates in epistemology*. Malden, MA: Blackwell, pp. 285-295.
- Langton, R. (1992) Duty and desolation. *Philosophy* 67, 481-505.
- Lewis, C.S. (1960) *The Four Loves*. London: Geoffrey Bles.
- Lintott, S. (2015). Friendship and bias: ethical and epistemic considerations. *Journal of Social Philosophy* 46(3), 318-339.
- Malcolm, F. (2018a). How to insult and compliment a testifier. *Episteme* 15 (1): 50-64.
- Malcolm, F. (2018b). Testimonial insult: a moral reason for belief? *Logos and Episteme* (1), 27-48.
- Marušić, B. (2017). Trust, reliance and the participant stance. *Philosophers' Imprint* 17(17), 1-10.
- Marušić, B. & White, S. (2018). How can beliefs wrong? A Strawsonian epistemology. *Philosophical Topics* 46(1), 97-114
- Mason, M. (2010). On shamelessness. *Philosophical Papers* 39(3), 401-425.

- Matthes, E. H. (2016). Love in spite of. *Oxford Studies in Normative Ethics* 6, 246-262.
- McHugh, C. (2015). The illusion of exclusivity. *European Journal of Philosophy* 23(4), 117-1136.
- Medina, J. (2012). *The epistemology of resistance: gender and racial oppression, epistemic injustice, and the social imagination*. New York, NY: Oxford University Press.
- Mill, J. S. (1863/2002). *Utilitarianism*. 2nd edn. G. Sher (ed.), Cambridge, MA: Hackett Publishing Co.
- Morrison, T. (2005). *Beloved*. London: Vintage Books.
- Murphy, M. C. (1999). The simple desire-fulfilment theory. *Noûs* 33(2), 247-272.
- Murphy, M. C. (2001). *Natural law and practical rationality*. New York, NY: Cambridge University Press.
- Nozick, R. (1974). *Anarchy, state, and utopia*. New York, NY: Basic Books.
- Oluo, I. (2015). I understand why people believe sexual predators before victims. I did. *The Guardian*. 22nd January 2015. Available at: <https://www.theguardian.com/commentisfree/2015/jan/22/why-people-believe-sexual-predators-before-victims> [Accessed on 30th January 2017.]
- Paul, S., and Morton, J. (2018). Believing in others. *Philosophical Topics* 46(1), 75-95
- Parfit, D. (1984). *Reasons and Persons*. Oxford: Clarendon Press.
- Pettit, P. (2004). Hope and its place in mind. *The Annals of the American Academy of Political and Social Science* 592, 152-165
- Piller, C. (2009). Desiring the Truth and Nothing but the Truth. *Noûs*, 43(2), 193-213.
- Piller, C. (2016). Evidentialism, transparency, and commitments. *Philosophical Issues* 26(1), 332-350.
- Preston-Roedder, R. (2013). Faith in humanity. *Philosophy and Phenomenological Research* 87(3), 664-687.
- Protasi, S. (2016). Loving people for who they are (even when they don't love you back). *European Journal of Philosophy* 24(1), 214-234
- Railton, P. (1986). Facts and values. *Philosophical Topics* 14, 5-31.
- Rawls, J. (1971). *A theory of justice*. Cambridge, MA: Harvard University Press
- Reed, C. (Director). (1949). *The third man*. [Film]. London: London Films UK.
- Reisner, A. (2009). The possibility of pragmatic reasons for belief and the wrong kind of reasons problem. *Philosophical Studies* 145(2), 257-272.
- Rinard, S. (2015). Against the new evidentialists. *Philosophical Issues* 25(1), 208-223.
- Rinard, S. (2017). No exception for belief. *Philosophy and Phenomenological Research* 94(1), 121-143
- Rinard, S. (forthcoming). Equal treatment for belief. *Philosophical Studies*, 1-28.

- Roeber, B. (2018). The pragmatic encroachment debate. *Noûs* 52(1), 171-195.
- Rosen, G. (2001). Nominalism, naturalism, epistemic relativism. *Philosophical Perspectives* 15, 69-91.
- Rorty, A. O. (1987). The historicity of psychological attitudes: love is not love which alters not when it alteration finds. *Midwest Studies in Philosophy* 10(1), 399-412.
- Ross, J. and Schroeder, M. (2014). Belief, credence, and pragmatic encroachment. *Philosophy and Phenomenological Research* 88(2), 259-288.
- Schlink, B. (1997). *The reader*. London: Phoenix House.
- Schoenfield, M. (2013). Permission to believe: why permissivism is true and what it tells us about irrelevant influences on belief. *Noûs* 47(1), 193-218.
- Schroeder, M. (2012a). Stakes, withholding, and pragmatic encroachment on knowledge. *Philosophical Studies* 160(2), 265-285.
- Schroeder, M. (2012b). The ubiquity of state-given reasons. *Ethics* 122(3), 457-488.
- Schroeder, M. (2013). State-given reasons: prevalent, if not ubiquitous. *Ethics* 124(1), 128-140.
- Schroeder, M. (2018a). When beliefs wrong. *Philosophical Topics* 46(1), 115-127.
- Schroeder, M. (2018b). Rational stability under pragmatic encroachment. *Episteme* 15(3), 297-312.
- Scruton, R. (1986). *Sexual desire: a moral philosophy of the erotic*. New York, NY: Free Press
- Shah, N. (2003). How truth governs belief. *Philosophical Review* 112(4), 447-482.
- Shakespeare, W. (1988). *Othello*. D. Bevington (ed.), USA: Bantam Books.
- Shakespeare, W. (2006). Sonnet 130. In S. Greenblatt (ed.), *The Norton anthology of English literature: volume one*. 8th Edn. New York and London: W. W. Norton & Company, p. 1074.
- Sidgwick, H. (1874/1962). *The methods of ethics*. London: Palgrave Macmillan.
- Siegel, S. (2017). *The rationality of perception*. Oxford: Oxford University Press.
- Stanley, J. (2005). *Knowledge and practical interests*. Oxford: Oxford University Press.
- Stanley, J. and Hawthorne, J. (2008). Knowledge and action. *Journal of Philosophy* 105(10), 571-590.
- Strawson, P. F. (1962). Freedom and resentment. In Gary Watson (ed.), *Proceedings of the British Academy, Volume 48: 1962*. Oxford: Oxford University Press. pp. 1-25.
- Stroud, S. (2006). Epistemic partiality in friendship. *Ethics* 116(3), 498-524.
- Stump, E. (2010). *Wandering in darkness: narrative and the problem of suffering*. Oxford: Oxford University Press.
- Taylor, G. (1976) Love. *Proceedings of the Aristotelian Society* 76, 147-164.
- Telfer, E. (1970). Friendship. *Proceedings of the Aristotelian Society* 71, 223-241.
- Thomas, L. (1987). Friendship. *Synthese* 72(2), 217-236.
- Thomas, L. (1993). Friendship and other loves. In N. K. Badhwar (ed.), *Friendship: a philosophical reader*. Ithaca, NY: Cornell University Press, pp. 48-64.

- Velleman, J. D. (1999). Love as a moral emotion. *Ethics* 109(2), 338-374.
- Velleman, J. D. (2001). The genesis of shame. *Philosophy and Public Affairs* 30(1), 27-52.
- Wallace, D. F. (2004). Oblivion. *Black Clock* 1, 121-145.
- White, Richard. (1999). Friendship and commitment. *The Journal of Value Inquiry* 33, 79-88.
- White, Roger. (2005). Epistemic permissiveness. *Philosophical Perspectives* 19(1), 445-459.
- White, Roger. (2013). Evidence cannot be permissive. In M. Steup, J. Turri, and E. Sosa (eds.), *Contemporary debates in epistemology*. Malden, MA: Wiley-Blackwell.
- Williams, B. (1981). Persons, character, and morality. In J. Rachels (ed.), *Moral Luck*. New York: Cambridge University Press.
- Williams, B. (2002). *Truth and truthfulness: an essay in genealogy*. USA: Princeton University Press.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.
- Wolf, S. (1992). Morality and partiality. *Philosophical Perspectives* 6, 243-259.