Beliefs That Wrong

by

Rima Basu

A dissertation submitted in partial fulfillment
of the requirements for the degree of Doctor of Philosophy (Philosophy)
at The University of Southern California
June 2018

Doctoral Committee:

Ralph Wedgwood, University of Southern California (Chair)
Mark Schroeder, University of Southern California
Stephen Finlay, University of Southern California
Kenny Easwaran, Texas A&M University
Aimee Bender, University of Southern California (External)

I dedicate this dissertation to my parents. I wish I could've been the daughter you deserved. To my wife, you are more than I deserve. And to my brother, you owe me.

ACKNOWLEDGEMENTS

I've failed, a lot. To say I did horribly on the GRE would be an understatement. The first time I applied to graduate school, I was rejected by every school. An end-of-year report after my first year at USC reads: "[at] this point in her career, it is not yet clear to me whether Rima is going to be able to be a successful student in our program." In my second year, I improved, but an end-of-year report reminds me that I would "need to continue on a steep curve of improvement." After I finally started a dissertation project, I ended up scrapping it and starting again on a new topic, and at some point I stopped counting rejections.

We don't often talk about our failures, and I know what you're thinking, an acknowledgements section seems an odd place to bring them up. But, when we don't acknowledge our failures, it can be easy to forget that these stories of failure and rejection are not unique. Everyone has these stories. And if not for the unwavering support of everyone who has guided me, supported me, and most importantly, believed in me, any one of these failures might have derailed me. I have been surrounded by more love and support than I deserve, and I have a lot of people to thank.

I want to start by thanking my professors at the University of Toronto Mississauga (UTM). When I chose my undergraduate institution, I didn't put much thought into it. UTM was a 10-minute drive from my parents' house; it was convenient. Little did I know that that undergraduate institution had one of the best collections of philosophy faculty in the world. I owe everyone there a lot. Amy Mullin was chair at the time and responsible in large part for the great undergraduate experience I had. Likewise, Jennifer Nagel, Marleen Rozemond, Bernard Katz, Jonathan Weisberg, Mohan Matthen, Diana Raffman, and Jacqueline Brunning all played a large role. In particular, I must say a special thanks to Sergio Tenenbaum, Paul Franks, Phil Clark, Gurpreet Rattan, and Andrew Sepielli. Not only did they write my letters for grad school, but they are also responsible in large part for the kind of philosopher I am today, as well as the kind of philosopher I hope to be.

I remember losing hours in Sergio's office talking about material we'd covered in class (and material we hadn't). Any Kantian sympathies that emerge in my work are because

of seeds that Sergio planted. I'm also grateful to both Sergio and Phil for being at my very first SLACRR and introducing this petrified 2nd-year grad student to famous people whose names she'd only ever seen on paper. Phil is also responsible for my love and admiration of Philippa Foot (whose portrait, painted by Renee Bolinger, hangs above my desk). Paul's mentorship and patience while I was an undergraduate fellow at the Jackman Humanities Institute is a model I can only hope to emulate one day. The paper I wrote for Andrew's class (after much revision and input from everyone at UTM) got me into graduate school. And Gurpreet was there for me with words of encouragement when I was at my lowest. Gurpreet convinced me that I should take that extra year Toronto was offering in the MA program and just get better and prove that I could do work at the graduate level. Second time applying to grad school things did go better, and I'm forever indebted to everyone at UTM and the University of Toronto generally for all the support and guidance they offered.

When I got to USC I entered firm in the knowledge of three things. First, I was going to work at the intersection of philosophy of language and metaethics. Second, I resolved to never write about reasons. Third, Los Angeles is the best city in the world. Any interest I had to continue working in philosophy of language evaporated after our first-year proseminar. And although I ended up writing about reasons (for a bit), my love for Los Angeles has never wavered. So, one outta three ain't bad.

There are many people at USC I have to thank. First, there is my cohort—Erik Encarnacion, Caleb Perl, and Steve Bero. In addition, Beth Snyder might as well have been part of our cohort. Beth's knack for storytelling is really quite something. If you ever have the opportunity to meet her, you should ask her about that time she got a haircut. Erik, Caleb, and Steve made me feel welcome. At first I was intimidated by them, but then Caleb got lost on the way back to his car, Erik confessed his fears, and Beth dished all the dirt on Steve. I'm grateful that they were my cohort and consider them the older siblings I never had. To this day I also remain jealous of Steve's salt and pepper hair.

In addition to my cohort, there was a wonderful group of older graduate students who immediately made me feel welcome and part of a community. I had only just arrived in Los Angeles when I was invited to come out for dinner and Shakespeare in the Park. To Matt Lutz, Ben Lennertz, Justin Snedegar and Emmy Feamster, Marina Folescu and Alex Radulescu, Justin and Amanda Dallmann, Johannes Schmitt, Jen Liderth, Aness and Myles Webster, Shyam Nair, Julia Staffel and Brian Talbot, Alida Liberman and Joshua Crabill, Ara Astorian, Kenneth Silver, Matthew Babb, Abelard Podgorski, Indrek Reiland, Sam Shpall, Lewis Powell, and Keith Hall: thank you. I've also benefited greatly from discussing my work with various post-docs that have visited USC, from

Bryan Roberts, Alex Sarch, Leif Hancox-Li, to Monica Solomon: thank you as well.

Now, as lovely as Erik, Caleb, Steve, Beth, the older graduate students and the post-docs were, there is am undeniable generational divide between us. None of them grew up playing Pokemon or arguing about who would play which Spice Girl during recess. In my second year, a new batch of first-years entered. There were now people with which I could debate the finer points of being a 90s kid. Weekly trivia nights with Mike and Amanda Ashfield, Jonathan Wright, Renee Bolinger, and Maegan Fairchild were what kept me going through the pressures of second year and the ever-looming second-year review process, which would determine if I was worthy enough to continue in USC's program. Every Thursday night we would meet at Barney's Beanery in Oldtown Pasadena, and we'd just unwind. And although we were really bad at trivia, we were really good at the free raffle that accompanied trivia. I cannot stress enough how important those nights were to me and how important those friendships became. But then, Mike and Amanda decided to have a kid, and I fell in love with a girl that lived in Santa Monica. Trivia nights weren't ever quite the same after that.

But before I move on from them, let me add that Mike and Amanda are two of the kindest people I know and they're raising two amazing kids. I also couldn't have asked for a better teaching partner than Mike. Mike approaches every topic he comes across as though it is inherently interesting, and that's a quality I greatly admire. Renee is stubborn and determined and has a single-mindedness that is admirable. Jonathan's got that gift for getting people out of their shell and making them feel welcome, and Maegan's not only extremely kindhearted, she's also got that talent of explaining your own research back to you in a way that's better than you ever could've stated yourself. Together, with my cohort, I couldn't have asked for a better group of people.

I also owe an additional debt of gratitude to Tanya Kostochka. Tanya is one of the few people who has ever witnessed truly cynical Rima. I am surrounded by a lot of optimists, a lot of folks who just want to believe the best of others. I am grateful for all of them, but I'm also grateful for Tanya as she's is also someone that can cut through the noise and see people for who they are and approaches interpersonal relationships with a healthy dose of cynicism. I also have to thank August Gorman and Nathan Howard for long conversations about philosophy and beds that I've crashed on. I also regret not getting to know Nicola Kemp, Jesse Wilson, and Christa Peterson better.

Another magical event that not only created deeper friendships, but also resulted in the formation of many dissertations, is John Hawthorne's seminar. The seminar was supposed to be about publishing. We were supposed to bring to that class an idea or a paper that we hoped to turn into a publishable paper. Each week we would meet, present,

give feedback, and then get drinks. And over the course of those fifteen weeks, we each discovered that the small paper idea we had, well, to do it justice it'd have to be more than one paper. That class is responsible for generating at least four dissertations at USC, and I'm forever grateful to John Hawthorne and my fellow classmates Tanya Kostochka, Mike Ashfield, Maegan Fairchild, Renee Bolinger, Matt Babb, and Matt Leonard for the feedback they offered over that semester.

It is not only USC graduate students I have to thank, but also UCLA graduate students. From hosting me as a prospective to long conversations about the nature of belief, I am grateful to call Greg Antill a friend. In addition, talking philosophy with Laura Gillespie (also the best wedding officiant Gabby and I could've asked for), David Friedell, Kim Johnston, Eric Tracy, Will Reckner, Sarah Beach, Ally Peabody, and Olufemi Taiwo has always been a delight. I guess I should also thank Kevin Lande, but that's more so for being a real person and not making me feel like we have to talk philosophy all the time.

In addition, there is The Compound: a set of duplexes right on the edge of West LA and Santa Monica where Gabbrielle Johnson, Gabriel Dupre, Ayana Samuel, Amber Kavka-Warren, Eleonore Neufeld, Andrew Stewart, Noah Gordan, Jasmine Gunkel, and I live. And where Isaac Khalaf and Laura Heida used to live. The conversations I've had with everyone at the compound, and the friendships that have been made over the years, have been an important part of my growth as a person and a philosopher, and to all my compound roommates: I am forever grateful. In particular, Gabe is secretly a big softie that genuinely cares about others and would do whatever he needed to do to make things work out for them. Amber holds herself to too high a standard and will succeed at anything she sets her mind to (except parking). Ayana is everyone's loveable grump, and Elli her foil.

Finally, we get to my committee. The friends we choose and the people we surround ourselves with are a reflection of the kind of person we are and the kind of person we hope to be. And with that thought in mind, I found myself in the lucky position of being able to work with not only excellent philosophers, but philosophers of excellent character. I cannot stress enough that I made a very conscious decision with regard to who I chose to work with. I chose each with the following thought in mind: if one day someone were to say that I remind them of any member of my committee I would not have to qualify or feel ashamed, e.g., "although x would never read anything...". Each member of my committee sets a standard for mentorship I not only look up to and hope to emulate, but they're also genuinely good people.

In Ralph Wedgwood, one finds patience and a deep well of intellectual curiousity. Talking with Ralph can be a lot like researching a topic on Wikipedia. You start with

rational requirements and you end up discussing the Agrarian Revolution of the mid-17th century (or the philosophical equivalent). Ralph joined USC in my second year with Ralph I found the freedom to explore whatever topics I was interested in and a patient conversational partner. That freedom was crucial to me finally landing upon this topic that has consumed me, and that I expect will continue to consume me. It can be easy to develop a narrow mind when working on narrow topics, but Ralph will always remind you of just how much larger everything is. I cannot say enough nice things about Ralph and I look forward to a long friendship.

In Mark Schroeder, one finds structure and foresight. Mark was one of the key reasons I chose USC. When I visited USC as a prospective student, we talked for an hour in his office and then for over an hour while crawling along the 10. My writing sample had also been on expressivism and epistemic modals, so of course it made sense that I wanted to work with Mark. However, my interest in expressivism waned and I did so horribly in my first class with Mark that I was afraid to keep working with him because surely he thought I was irredeemably stupid. The two quotes at the start of the acknowledgements are from end of year reports written by Mark. But, for some reason, Mark continued to believe in me. He continued to ask to read my writing and provide feedback, and once I was finally able to articulate the problem of the supposedly rational racist, we both were hooked. None of this yet speaks to the qualities I mentioned at the start of the paragraph: structure and foresight. Whereas with Ralph I had the freedom to just explore, Mark is who I turned to when I needed to really nail down an idea. Mark can't see the future, but he can see exactly where an argument leads. It's not a skill I've figured out yet, but until I do, I know that I'll always still be able to send Mark any thoughts I have.

In Kenny Easwaran, one finds one of the most genuinely good and kind person to have ever existed. I do not think there is a bad bone in his body and very few people have come close to being as genuine (only other Canadians come to mind). I thought I hated epistemology until I took a class with Kenny and found out that epistemology wasn't just about knowledge, but also about belief, the degrees those beliefs come in, justification, evidence, etc. Kenny opened up a whole area of philosophy to me, but most importantly I have never left a conversation with Kenny feeling as though I can't do philosophy, no matter how hard the topic. Finally, not only does Kenny have the gift of being able to explain complex topics in a simple way, he and Ralph share the following characteristic: they know a lot about a lot of stuff that's not philosophy. As a result, I've learnt not only a lot about philosophy from Kenny, but also about city planning and bike lanes. Losing Kenny to Texas A&M was a huge loss for USC, but I'm lucky to have been able to continue to work with him.

In Steve Finlay, one finds an uncanny ability to inhabit differing perspectives when reading a piece of work to figure out all the ways in which a sentence is likely to be misread. This careful attention and broad perspective is one that I hope to be able to offer for my own students one day. Steve is also part of the camp that recognizes that to do philosophy, one must read more than just philosophy. I'm also grateful for our discussions about Doctor Who. On his advice I decided to give the show a second chance (after becoming frustrated with Stephen Moffat), and it turns out that Season 10 is actually quite good. Thanks, Steve!

Due to "understood" rules that weren't explicitly listed, Pamela Hieronymi was not able to be an official member of my committee. She did, however, play the part in every way. In Pamela Hieronymi, one finds a standard of precision and rigour to aspire to (and that is characteristic of UCLA's reputation—Gabby made me add that). Pamela is someone who understands that there is something we have to get right when we're doing philosophy. We're not just asking questions and clumsily positing answers, there's hard work to be done and not everyone's up to the task. Over the course of writing this dissertation, Pamela has been my hardest critic, and for that I am deeply grateful. Pamela has provided that useful outsiders perspective. If I could convince Pamela of a point, then I was golden. As to whether I've succeeded... luckily I've got a few more years in the profession to keep trying.

Altogether, I couldn't have asked for a better committee. There are also the outside members I should thank. I'm grateful to Roumyana Pancheva for playing the role of the outside member for my qualifying defense, and to Aimee Bender for stepping in on such short notice for the dissertation defense.

There are also faculty that I didn't work closely with, but who've contributed to any successes I've had at USC: Shieva Kleinschmidt, Robin Jeshion, Janet Levin, and Gabriel Uzquiano. In particular, I want to say a bit more about both Shieva and Robin. Shieva provided friendship, guidance, and a buddy to watch superhero movies with. I regret the way in which we drifted apart once I moved to west of the 405, but I hope that there will be more ice cream and Marvel in our future. In addition, Robin Jeshion is the best teaching mentor—and mentor in the profession more generally—I could have asked for.

In addition, the office staff during my time at USC have been absolutely central to any successes I've had: Cynthia Lugo, Corey Resnick, Barrington Smith-Seetachitt, J.N. Nikolai, Natalie Schaad, and Rujuta Parikh. In particular, Cynthia, Barrington, and Natalie have saved me from many crises, and Natalie and Rujuta were the steady calming voice of reason and competence I needed during the job market. I can't say enough kind things about the amazing office staff that we've been blessed with. And although I haven't

had a chance to interact with the newest member of our office staff yet, Michele Root did process a reimbursement for me in less than an hour so I know she's going to be fantastic and we're lucky to have her.

Further, parts of this dissertation have been presented at the New Insights and Directions for Religious Epistemology Workshop at Oxford University, the University of Toronto Mississauga, Athena in Action at Princeton University and the Princeton Workshop in Social Philosophy, the Penn-Rutgers-Princeton Social Epistemology Conference at the University of Pennsylvania, the 2017 Central Division Meeting of the American Philosophical Association, Duke University, the University of British Columbia, Wayne State University, McMaster University, Georgetown University, Claremont McKenna College, Queen's University, Simon Fraser University, University of Missouri, and the 2018 Pacific Division Meeting of the American Philosophical Association. I've benefitted tremendously from feedback that I've received every time I've presented my work, and I have tried to thank individuals for their contributions to my thinking on these topics in footnotes throughout the dissertation. Inevitably I'll have missed someone, for that I am deeply sorry.

And finally, I've saved the two most important people for last: Gabbrielle Johnson and Kripke (the dog). I don't know where I'd be if not for Gabby. Gabby has been a constant source of love, support, friendship, guidance, every virtuous trait you can name. Gabby has been my copy editor, my sounding board, someone I can run any idea by and in just talking it through she'll make it better. Gabby is such a demanding interlocutor that for the sake of our relationship we tend not to talk philosophy, but I will never send something off until she's given it the once-over. She has made me a better person and a better philosopher, and I cannot thank her enough. I look forward to our living together for a hundred autumns.

Last, but not least, Kripke is the best dog that anyone could ask for. She is truthfully the only reason I work as hard as I do (when I do). One day I hope to be the dog-parent she sees in me already, and to return the favour of her many years of friendship with the giant yard she's always deserved.

TABLE OF CONTENTS

INTRODUCTION

The dissertation begins with the example of Spencer, the supposedly rational racist. Spencer argues that although it might be "unpopular" or "politically incorrect" to say this, the evidence supports believing that the black diner in his section will tip poorly. He insists that the facts don't lie. *The facts aren't racist*. If you were to deny his claim and believe otherwise, he'd challenge that it is *you* who engages in wishful thinking. It is *you* who believes against the evidence. You, not Spencer, are epistemically irrational.

An unfortunate fact (amongst many) is that we live in a society that has been shaped by racist attitudes and institutions. Given the effects of structural racism, Spencer's belief could be supported by the evidence. Spencer notes that it might make him unpopular, but he cares about the truth and he is willing to believe the unpopular thing. But, Spencer's belief seems racist. Spencer asks, however, how could his belief be racist if his beliefs reflect reality and are rationally justified? Moreover, how could he wrong anyone by believing what he epistemically ought to believe given the evidence?

The overarching goal of the dissertation is to give an account of the moral-epistemic norms governing belief that will help us answer Spencer and the challenge he poses. In answer, I argue that beliefs can wrong. At least one of the distinctive wrongs committed by a racist—and of central importance to this dissertation, the wrong committed by Spencer—plausibly lies in what they believe about another human being.

To bolster support for this intuitive thought, there are many cases we can turn to as *prima facie* evidence that we can wrong someone by what we believe of them. For example, one common formulation of the Christian Eucharistic confession, "we have sinned against you in thought, word, and deed", appeals to the idea that we can sin against God in *thought*, as well as in word and in deed. When loved ones believe the worst of us, it is tempting to think that we can demand an apology for the beliefs they hold, and not just their actions. Consider, for example, how common it is to hear the phrase, "You shouldn't have believed that of me!". Further, many people also think that we can wrong not only the living but also the dead when we believe the worst of them. As I will show in this dissertation, the idea that we can wrong someone by what we believe reveals itself in many places. But, you might counter, perhaps we simply mistaken when we talk this way.

I grant that this idea that beliefs can wrong is philosophically puzzling. The norms that properly govern belief are plausibly epistemic norms such as truth, accuracy, and evidence. Moral and prudential norms are taken to play no role in settling the question of whether to believe *p*, and they are irrelevant to answering the question of what you should believe. It would follow, then, that the person who believes of Barack Obama that he's more likely to be a valet than the President of the United States need only attend to the relevant statistical likelihoods before settling their belief. As we will see argued by the antagonist of this dissertation, Spencer, that they commit an injustice to Barack Obama when they mistake him for a valet ought to be irrelevant in assessing the belief itself (as opposed to, say, its consequences). This leaves us with several puzzling questions: can we wrong one another by virtue of what we believe about each other, i.e., can beliefs wrong? If not, why not? What then is going on in these cases? If so, how? What then does this tell us about the content and nature of the norms governing belief?

In this dissertation I explore this perceived conflict between the intuitive idea that we can wrong one another by what we believe about each other, and the philosophically orthodox idea that the only norms that relate directly to belief concern truth-related factors such as evidence. Motivated by the recognition that our epistemic practices exist in an unjust and non-ideal world, I argue that there are moral and social constraints on our epistemic practices. Specifically beliefs can wrong and the epistemic justification of our beliefs can be affected by the moral demands of our environment. That is, when it comes to what we should believe, morality is not voiceless: there is a moral dimension to the way in which our beliefs are justified. This thesis has important consequences for not only how we think about moral questions, e.g., what we owe to each other, but also how we think about epistemic rationality. My dissertation is thus divided into two parts. In the first, I make the case for the thesis that beliefs can wrong and I explain why, when, and who our beliefs can wrong. Then, in the second half I show how these moral considerations affect the epistemic rationality of Spencer's belief by developing an account of *moral encroachment*.

Proceeding chapter by chapter, the task of Chapter 1 is to explain why there's a wrong that should be located in belief. Many try to explain away the appearance that beliefs can wrong by locating the wrong elsewhere, e.g., in the actions of the agent. However, these alternative accounts only give us a *diagnosis* of Spencer. These accounts fail to both engage with Spencer and fully explain the moral wrong in these cases. To give a diagnosis that also *speaks to* Spencer, we must appeal to the surprising thesis that people can wrong other people in virtue of what they believe about each other, and not just in virtue of what they do. The task of Chapter 1, in short, is to make a case for *why we should think* there is a

wrong that is located in what an agent believes of another person as opposed to what they say or do.

In Chapter 2 I show that the phenomenon of wronging beliefs is more general than just Spencer's racially charged beliefs. Although you might not be harmed when someone believes something of you on the basis of statistical evidence, you are nonetheless wronged because a moral demand has not been met: the demand to be related to in a way that is characteristic of being an agent rather than an object. Despite the fact this thesis seems to make an implausibly wide range of things morally problematic, I suggest we switch from considering Spencer as our central example to considering what it feels like for the agent who is mistaken for a waiter on the basis of their race or presumed to be a bad tipper on the basis of their race. In many contexts if your melanin levels more closely match that of the waitstaff than that of the diners, it seems the belief that you are a waiter is supported by the weight of the evidence. Granting that there is nothing wrong with being a waiter, why then does it hurt when someone mistakes you for a waiter in this way? In this chapter I argue that this way of forming beliefs about others fails to meet the moral demand of recognizing others as agents, rather than objects. We fail, in short, in what we epistemically owe to each other.

We now arrive at a new problem. Although the evidentially rational belief might be morally wrong, the belief still seems like it might be *epistemically* fine. This suggests the following dilemma: given the evidence, you ought believe; but, given the moral considerations, you shouldn't. After all, in many real-world scenarios, Spencer seems right. That a diner is black can provide a substantial amount of evidential support for the belief that the diner will not tip well. If Chapter 1 and 2 are correct, it is morally impermissible to believe that he will not tip well. On the other hand, the belief seems epistemically justified. If one had more hands, one might go further and add that it would be epistemically irrational not to believe that he won't tip well. To refrain from believing seems to involve a wilful disregard of the evidence. So, what should you believe?

Many conclude that in such cases we face an irresolvable conflict between what morality requires and what we epistemically ought to believe. I disagree. In Chapter 3, I argue that we can resolve this conflict by appealing to a form of moral encroachment. That is, the epistemic justification of our beliefs can be determined, in part, by the moral demands of our situation, i.e., the moral stakes. As an upshot of this account, the epistemic considerations one weighs when trying to settle the question of what to believe include moral considerations. That is, moral and epistemic considerations need not stand at odds with one another. I argue that we must rethink the assumption that gives rise to these conflicts: the evidentialist assumption that only evidence matters for justification. Evidence

alone is not sufficient. I argue that non-evidential considerations—in particular, the moral stakes of a belief—must enter into the picture here. This is the phenomenon of *moral encroachment*. What is sufficient evidence does not supervene on what your evidence is. The moral stakes can affect whether a belief is justified.

Finally, in Chapter 4, I canvas reasons for dissatisfaction with the traditional answers to this supposed dilemma between moral and epistemic demands, e.g., that there is an all-things-considered ought that adjudicates conflicts. Given that the epistemic justification of our beliefs can be affected by the moral demands of our situation, I argue that we can prescind from questions about whether there is an all-things-considered ought that adjudicates the conflict. Spencer's belief is both morally objectionable and epistemically irrational. Further, it is a mistake to think that our moral lives and our epistemic lives are separate realms governed by conflicting obligations. As a final take away, to be moral we must not only do better, we must believe better as well.

CHAPTER I

The Wrongs of Racist Beliefs

We care not only about how people treat us, but also what they believe of us. Suppose I believe that you're a bad tipper given your race. Have I wronged you? What if you are a bad tipper? It is commonly argued that the way racist beliefs wrong is that the racist believer either misrepresents reality, organizes facts in a misleading way that distorts the truth, or engages in fallacious reasoning. In this opening chapter I begin by presenting a case that challenges this orthodoxy: the case of the *supposedly rational* racist. We live in a world that has been, and continues to be, structured by racist attitudes and institutions. As a result, the evidence might be stacked in favour of racist beliefs. But, if there are racist beliefs that reflect reality and are rationally justified, what could be wrong with them? Moreover, how do I wrong you by believing what I epistemically ought believe given the evidence?

To address this challenge, I argue that we must recognize that there are not only epistemic norms governing belief, but moral ones as well. This view, however, is at odds with the assumption that moral obligation requires a kind of voluntary control that we lack with regard to our beliefs. This background assumption motivates many philosophers to try to explain away the appearance that beliefs can wrong by locating the wrong elsewhere, e.g., in an agent's action. Further, even accounts that accept the thesis that racist beliefs can wrong restrict the class of *beliefs that wrong* to beliefs that are either false or the result of *hot irrationality*, e.g., the racist belief is a result of ill-will. In this chapter I argue that although these accounts will capture many of the wrongs associated with racist beliefs, they will at best be only *partial* explanations because they cannot explain the wrong committed by the supposedly rational racist.

The challenge posed by the supposedly rational racist concerns our epistemic practices in a non-ideal world. The world is an unjust place, and there may be many morally objectionable beliefs it justifies. I argue that to address the challenge posed by the supposedly rational racist we must seriously consider the thesis that people wrong others in virtue of

what they believe about them, and not just in virtue of what they do. In short, my goal in this chapter is to show that we miss *something* when we ignore the possibility of doxastic wronging. As for what that *something* is, that I will explain in Chapter 2.

1.1 The Motivating Challenge

1.1.1 The Supposedly Rational Racist

It is not up for debate that we live in a world that has been shaped by, and continues to be shaped by, racist attitudes and racist institutions. From the transatlantic slave trade, to anti-miscegenation laws, lynchings, redlining, and voter identification laws that "target African-Americans with almost surgical precision", racism is an unfortunate part of the fabric of our world.[1] It should not be surprising, then, that as a result of structural racism, there may be morally objectionable beliefs that are well-supported by the evidence. Further, some of the morally objectionable beliefs could be paradigmatic examples of *racist* beliefs. Although my focus in this chapter are racist beliefs and the ways they can wrong, what I say can also be easily extended to cover the possibility of rationally held *sexist* beliefs, *homophobic* beliefs, and other morally objectionable beliefs of this kind. The world is an unjust place and there may be many morally objectionable beliefs that it justifies.

Consider, for example, Miranda Fricker (2007, pp. 88)'s description of the following case. In *The Talented Mr. Ripley*, the character Marge is constantly being construed as hysterical by the males in the novel. Their dismissive attitude towards her and her testimony builds up to her eventual breakdown and having to be physically restrained. She becomes what she has been constructed as: a hysterical female that expresses herself in semi-contradictions who disregards the facts and is unable to keep a grip on her emotions. This fictional case illustrates how it's not unreasonable to think that if women are constantly told that they are hysterical, then they might start behaving in that manner. As a result, the sexist belief that women are hysterical may become well-supported by the evidence.[2]

Continuing this theme, let us consider the following case that is the motivating puzzle of this dissertation.

> *The Supposedly Rational Racist*. You shouldn't have done it. But you did. You scrolled down to the comments section of an article concerning the state

[1] Regarding North Carolina's discriminatory voter ID law, see: https://www.nytimes.com/2017/05/15/us/politics/voter-id-laws-supreme-court-north-carolina.html?

[2] For more on this point see Haslanger (2012) who draws on Hacking (1995, 1999)'s work on the "looping effects" of social kinds and structures. These social kinds can be self-fulfilling and self-reinforcing. Thanks to Fran Fairbairn and Jessie Munton for bringing this to my attention.

> of race relations in America, and you are now reading the comments. The comments on such articles tend to be predictable, but there is one comment that catches your eye. Amongst the slurs, the get-rich-quick schemes, and the threats of physical violence, there is the following comment: "Although it might be 'unpopular' or 'politically incorrect' to say this, I'm tired of constantly being called a racist just because I *rightly* believe of the black diners seated in my section that they will tip worse than white diners in my section." The user posting the comment, Spencer, argues that the facts don't lie, and he helpfully reproduces those facts. For example, he links to studies that show that on average black diners tip substantially less than white diners. The facts, he insists, aren't racist. If you were to deny his claims and were to believe otherwise, it would be *you* who is engaging in wishful thinking. It would be *you* who believes against the evidence. It would be *you*, not Spencer, who is epistemically irrational.

If we could dismiss Spencer's comment on the grounds that he is making a mistake, there would be no need to continue reading. The challenge posed by Spencer is that we cannot so easily dismiss his comment. From a study by Michael Lynn (2006), there is evidence of racial disparities in tipping practices that support Spencer's belief that black diners tip worse than white diners. Perhaps Spencer is, as he suggests, rationally believing in accordance with the evidence. Although his belief might *seem* racist and as such irrational, perhaps he is not guilty of making any kind of epistemic fault. That is, perhaps he is not misinterpreting the evidence or cherry-picking statistics to justify his pre-held prejudices. Perhaps the belief isn't the result of ill-will or hostility. Perhaps Spencer is right in his claim that it is not his fault that as a result of structural racism the evidence is stacked in favour of racist beliefs. Given that our world is racist, some of our beliefs will also be racist.

A rational racist, however, sounds like an oxymoron. Racism is often regarded as definitionally involving standard epistemic flaws such as false beliefs, obscuring or concealing or disregarding relevant information, organizing facts in a misleading way, and/or engaging in fallacious reasoning. For example, as Tommie Shelby (2016, pp. 22) argues, "in its most basic form, racism is an *ideology*: a widely held set of associated beliefs and implicit judgments that misrepresent significant social realities and function, through this distortion, to bring about or perpetuate unjust social relations."[3] The problem thus posed by the supposedly rational racist is that Spencer does not seem to be committing one of these

[3]This is a commonly shared view. See for example, Appiah (1990), Ikuenobe (2011), Lengbeyer (2004), Clough and Loges (2008), etc. I will explore similar accounts in more detail in Section 1.4.

standard epistemic flaws, but his belief nonetheless seems racist, and by believing what he does—i.e., that the black individual in his section will tip poorly—he seems to commit a wrong. But, if he is not guilty of an epistemic flaw, then it seems that we are at a loss to explain what wrong he commits.

Perhaps the reason we are at a loss to explain what wrong Spencer commits is simple: Spencer is not committing a wrong. Sometimes there are unpleasant things that we ought, at the very least, are permitted, to believe. For example, it might be unpleasant for me to believe that one day soon my favourite television show will come to an end (as many of my past favourite television shows have done, RIP Battlestar Galactica); nonetheless, it is something I must accept and it would be irrational for me to believe otherwise. This is precisely the line that Spencer advocates: there is at least some value to having accurate beliefs, even when those beliefs are unpleasant. The beliefs that Spencer holds, however, are not merely unpleasant. The beliefs that Spencer holds *seem racist*. But, again, one might counter that although the beliefs *seem* racist, it would be wrong to call them racist given that they reflect reality.

The challenge, then, is to explain why we should call Spencer's belief racist if his beliefs reflect reality and seem to be rationally justified. Perhaps by calling such a belief racist we are contributing to a conceptual inflation of the term and thereby making it less powerful for calling out things that are *actually* racist. Perhaps it is not racist to believe a black diner won't tip well. This is the challenge posed by the supposedly rational racist.

1.1.2 An Assumption and A Constraint

This is a troubling conclusion. I am not as willing as Spencer to give up on the compelling idea that there is something wrong with what he believes even when the belief might accurately reflect reality. Assuming the reader shares this hesitancy, this leaves us with the task of figuring out what Spencer does wrong.

As we move forward with this task, there is one assumption and one constraint I ask the reader to take on. First, I will be assuming that by believing what he does, Spencer commits a wrong.If we take it as given that Spencer commits a wrong, our task becomes explaining what wrong Spencer commits. To reject this assumption is to reject that there is a challenge posed by Spencer. That is, if Spencer does not commit a wrong, there is no puzzle that needs to be explained. The case of the supposedly rational racist, however, is puzzling because of how compelling this assumption is, i.e., that Spencer commits a wrong by believing what he does.[4]

[4]In this chapter I will argue that Spencer commits a wrong and the wrong is properly located in what he believes, and then in Chapter 2 I will give an account of *why* Spencer's belief wrongs.

Further, I also ask the reader to take on the following constraint on the range of acceptable answers to the case of the supposedly rational racist: when giving an explanation of the wrong committed by Spencer, we should, as much as possible, do our best to grant Spencer his conception of himself.[5] Spencer is not merely a philosopher's concoction to counterexample or test the limits of a proposed theory. Spencer exists. Spencer might be our brother sharing infographics of "black-on-black crime"; Spencer might be our aunt that everyone politely tolerates at Thanksgiving dinner when she complains about immigration reform given her belief about the correlation between undocumented immigrants and crime rates. Your aunt is always careful to choose her words so as to avoid the misconception that she's being racist, to present herself as though she is simply compelled by the facts. We could easily write off our brother or our aunt as making some kind of mistake, but just as I suggest we try to engage with Spencer and not believe the worst of him, we certainly do not want to believe the worst of our loved ones.

We could tell our brother that the fact that he's constantly sharing infographics about black-on-black crime reveals his motivated irrationality, e.g., his ill-will towards black people. His belief that 90% of African-American homicides are committed by African-Americans might not be false, but that he claims to be justified in believing that his black neighbour is dangerous is the result of both a kind of *motivated irrationality* and a misunderstanding of statistics. That is, our brother is clearly willing to believe anything about his black neighbour as long as it is negative. We could also say to our aunt that her belief about a correlation between crime rates and undocumented immigrants goes in the opposite direction, i.e., that undocumented immigrants commit less crime than American citizens. I grant that these descriptions are often accurate, and that encountering someone who truly is as Spencer envisions himself is rare. Nonetheless, the case of Spencer stands apart from our brother and our aunt. I ask that we take on the following constraint regarding our diagnosis of Spencer: we should try to describe him in a way that he would accept. Ideally, we wish not only to explain the wrong/s that Spencer commits in believing what he does, we also do not want to rule out the possibility of *engaging* with Spencer. Engagement, however, requires working with this constraint and not speculating about the nature of Spencer's psychology or making claims about various character

[5] You might worry here that this constraint conflicts with the assumption that Spencer commits a wrong. After all, surely Spencer does not accept that he commits a wrong. Hence, I only claim that we should, *as much as possible*, do our best to grant Spencer his conception of himself. Crucial to Spencer's self-conception is that he is epistemically rational. To show my hand, I plan to show over the course of this dissertation that it is not at odds with epistemic rationality to accept that others can place demands upon on how we respond to evidence. If Spencer grants that for his friends, then it is a short step to granting that for vulnerable minorities like the black man in his section. It is not, I argue, at odds with the demands of epistemic rationality to exhibit epistemic partiality.

flaws he might exhibit. To engage, ideally we should strive for a response to Spencer that would show that we can grant his conception of himself and still, by his own lights, he should not believe what he believes.

1.1.3 The Shape of Things to Come

Returning to the task of explaining how Spencer commits a wrong by believing what he does, the first problem we face is explaining why Spencer's belief should be called 'racist'. The first problem is that the belief seems to accurately reflect the evidence. Second, perhaps the term 'racism' should be reserved for morally problematic *false* beliefs. Further, when it comes to what the term 'racism' should be reserved for, perhaps it shouldn't apply to beliefs at all. Perhaps the term should be reserved for the actions that follow from a belief.

In Section 1.2 I consider this strategy of locating the wrong *downstream* from the belief. According to this strategy, Spencer ought believe what he does, but where he would go wrong is *in acting* on the belief. Against this proposal, I argue that if we restrict ourselves to locating the objectionability of racist beliefs in the downstream effects of the beliefs, we fail to capture a number of other ways in which racist beliefs can wrong. We care not only what people do to us, but also about the attitudes and beliefs they hold about us. This strategy, then, leaves us with only a partial explanation of the wrong of racist beliefs. For a more complete explanation, we must recognize that there are other places the wrong can be located.

In Section 1.3 I consider accounts that locate the wrong of racist beliefs *upstream* from the belief, i.e., the many processes that contribute to the agent forming the belief. For example, the wrong might lie in an agent's cognitive limitations or predispositions. The problem with these accounts is twofold. First, they violate the working constraint that we grant Spencer his conception of himself; second, the accounts also run into trouble with the assumption that deontic notions do not apply to beliefs because we lack sufficient voluntary control over our beliefs. This second assumption is part of the motivation of locating the wrong elsewhere. Although you could avoid endorsing the thesis that beliefs can wrong by locating the wrong upstream from belief, what little control we exercise over what we believe is surely more than the control we exercise over any upstream feature of our cognitive architecture. The irony for upstream accounts is that the motivation for locating the wrong upstream is also an objection against attempts to locate the wrong upstream. If beliefs can't wrong because we lack sufficient control over beliefs, then it's not clear how cognitive predispositions could wrong either. I conclude this section by noting that we could nonetheless say that Spencer's wrong is that he ignores the risk of

creating a biased cognitive system when he believes that his black diner, Jamal, will tip less well than his white diner, James. However, this requires endorsing the thesis that beliefs can wrong.

In Section 1.4 I turn to accounts that locate the wrong in what an agent believes. I call these accounts *restricted* accounts of the wrong because they specify two conditions for racist beliefs that wrong. First, the belief is false, i.e., *the false condition*. Second, the belief is a reflection of a character flaw such as ill will or racial disregard, i.e., *the hot condition*. Spencer does not satisfy the first condition—his belief *seems* rationally justified and it is not false. Further, the second condition requires denying the working constraint that we grant Spencer his conception of himself. I grant in the vast majority of cases people who claim to be like Spencer satisfy either the false condition or the hot condition. However, the case of Spencer illustrates that they needn't. Thus, these accounts are incomplete.

This brings us to the final section. In Section 1.5, I defend the intuitive plausibility of the thesis that beliefs can wrong. As we will have seen in the previous sections, the standard explanations for the wrong committed by Spencer do not work. I suggest, then, we rethink the assumptions that both stand in the way of the thesis that beliefs can wrong, and prevent us from positing moral-epistemic norms governing belief. Doing so allows us to preserve the working assumption and working constraint, i.e., that by believing what he does Spencer commits a wrong, and in explaining the wrong we should describe Spencer in a way that he would accept. As I have been noting, however, the thesis that beliefs can wrong conflicts with a commonly accepted background assumption regarding the nature of belief and moral disapprobation: beliefs lack the kind of voluntary control needed for our beliefs to *wrong*. My argument against this final challenge is indirect: having ruled out all plausible alternatives for locating the wrong committed by Spencer, we arrive at the conclusion that it is Spencer's belief itself that wrongs. The final consideration, then, is a choice between Spencer not wronging at all or Spencer having wronged through what he believes. It seems intuitively obvious to me that the supposedly rational racist does wrong. And as I shall show in this final section, this is more obvious, in fact, than the theoretical normative posit that voluntary control is always required to wrong. To give Spencer a pass on the basis of a theoretical technicality does a disservice to the common folk-intuition that something has gone wrong in the case of Spencer. It is thus worth considering how it could be true that beliefs can wrong.

1.2 Beliefs Aren't Racist: Locating the Wrong Downstream from Belief

Our task is to explain the wrong Spencer commits. I have already showed my hand by noting that the answer I find most compelling is that by *believing* what he does Spencer does something wrong. This is not, however, the most obvious answer. To get to this answer, we must first work our way through other accounts that have been offered and show their shortcomings.

In this section, my focus is on *downstream* accounts of the wrong. According to downstream accounts of the wrong, it is not the *belief* itself that wrongs: rather, it is the downstream effects of the beliefs that wrong, e.g., the agent's actions. Spencer only does something wrong given what he believes insofar as he *acts* on what he believes. Here is the intuitive thought behind such accounts: we have been assuming that what Spencer believes is an accurate reflection of the evidence. Now we can ask, how might Spencer wrong a black diner by simply *believing* that they are going to tip less well than a white diner? Surely he doesn't wrong the black diner unless he *treats* the black diner differently than the white diner. For example, it is clear that he would wrong the black diner if he were to provide worse service on the basis of the diner's race. But, if he simply believes? What's the harm (or wrong) in that? If no racist actions follow from the belief, then he does not do anything wrong. The downstream theorist, thus, can accept a restricted version of my claim that people wrong each other in virtue of what they believe about each other. People wrong each other, this theorist claims, in virtue of what they believe about each other, but only insofar as how they *treat* each other is a reflection of what they believe about each other.

As we can see in this formulation, it is implicitly assumed that our beliefs toward one another are beyond the purview of what's regarded as proper or improper treatment. There are two reasons for this restriction. The first is an assumption about control and responsibility: moral obligation requires a kind of voluntary control that we lack with regard to our beliefs. If you accept this assumption, then you would be driven to locate the wrong downstream from belief in something that we do have voluntary control over, i.e., what we say and what we do. The second reason stems from a worry about the conceptual inflation of the term 'racism'. If you are concerned that the terms 'racism' and 'racist' are being over-applied and losing their force as terms of moral reproach, you might be driven to reserve the term 'racism' and 'racist' for the worst race-based harms, and be disinclined to apply the terms to what an agent believes. Although distinct reasons, both reasons can be summed up with the following slogan: beliefs don't wrong people, people wrong

people. Alternatively, people wrong people through their words and through their deeds, not simply what they believe.

I turn now to considering the case for downstream accounts of the wrong in more detail. As I'll show, such an account will be too narrow in scope. It will fail to capture the more subtle, insidious, and more prevalent forms of racism. Further, there is intuitive reason to think that the demand for proper treatment extends to what people believe of us, not simply how they act towards us or what they say of us. Although this would put us at odds with the assumption about control and responsibility, there might be reason to think that that assumption may not be as strong as it is commonly assumed to be. I leave the argument for that point, however, to Section 1.5. Let us begin, then, with the concern about the conceptual inflation of the term of 'racism' and why that leads us to locate the wrong of racist beliefs solely in the actions of an agent.

A prominent example of an argument for why the wrong of racism is located only in action comes from J. Angelo Corlett (2005). Corlett argues that the only kinds of ethnic prejudice and discrimination that rise to the level of racism, what he calls "racism proper", are overt negative actions or inactions that are harmful and hurtful to victims. Further, these actions or inactions must be egregious enough to be punished by law. His grounds for arguing this come from the belief that "it does little good (though I suppose some good, on some occasions) to call something racist when the law cannot and should not effectively deal with it" (pp. 578-9).

There are two parts to Corlett's view that we need to keep separate. The first claim is that "racism proper" is located only in action, and the second is that actions and inactions that the law cannot and should not deal with are rightly excluded from "racism proper". The second claim is Corlett's grounds for the first. There are, however, good reasons for rejecting the second assumption. But first, why might we want to delineate a special case of racist wrongs as rising to the level of racism proper? For this argument, we must turn to Lawrence Blum (2002)'s influential distinction between, on the one hand, racist beliefs and racist actions and, on the other, and beliefs and actions that are merely rac*ial* in nature.

Blum argues that the term 'racism' has been the victim of a conceptual inflation. The term has been overused and, in serving as a term of moral reproach, it has joined other vices such as dishonesty, cruelty, hypocrisy, etc. Blum argues that we need a more varied and nuanced moral vocabulary in order to capture the sense in which not every instance of racial conflict, insensitivity, discomfort, miscommunication, exclusion, injustice, or ignorance is "racist". He argues that insensitivity, ignorance, discomfort, etc. in actions and beliefs are lesser ills, and the term 'racism' should be reserved for the worst race-based

wrongs, rather than used to categorize and condemn all race-related ills. Applying this account to our case of Spencer, Corlett and Blum would say that Spencer's belief might be morally bad in many ways, but it is not racist.

However, this policing of when the term 'racism' can be employed is problematic. Racism takes various forms and admits of degrees of wrong. Lynching, redlining, and the transatlantic slave trade are all particularly egregious examples of racism. Although being thought of as a bad tipper on the basis of your race is not as bad as not being allowed to marry outside of your race, it does not follow that the former is not racist. Blum makes the distinction because he is concerned with how dialogue about racism might be impeded by charges of "racism". And Corlett makes this distinction between racism and "racism proper" because he is concerned with legal remedies in a racist society. But, as Shelby (2003, pp. 125) argues, these approaches downplay the value of using the term.[6]

> Many people in the US now think that racism is largely a problem of the past and are thus little concerned with racial issues. The expansive conception of racism with its strong negative valence could encourage those who are complacent to be more vigilant and circumspect in the racial realm. Perhaps they should accept the comparatively small burdens of hyperbolic uses of "racism" as the cost of achieving racial justice and of publicly affirming the humanity of those who are racially stigmatized.
>
> [...]
>
> Moreover, though many of Blum's "non-racist" race-related ills are lesser moral evils taken singly, their cumulative and wide-ranging impact on the life-chances of subordinate "races" creates a heavy burden that is properly thought of as a form of oppression. Calling these lesser wrongs "racist" reminds us of the seriousness of these sources of disadvantage, even if their perpetrators are not moral monsters.

These considerations allow us to reject attempts to distinguish between different degrees of the wrong, e.g. racial wrong, racism, and racism proper, and argue that there is good reason for calling all the ills that fall on that spectrum racist or instances of racism. And although Corlett's argument for locating the wrong of racism solely in action or inaction rested on the assumption that there was a special class of wrong, i.e. racism proper, that could only be located in action or inaction, we have not yet undermined the intuitive

[6] Blum later changes his mind in Blum (2004) and argues that there is a plurality of racial ills that could all be regarded as racist.

case for thinking that the wrong of racism resides only in action. So, let me now turn to providing this intuitive case and then challenging it.

> *Ferguson*. Consider Officers Stella and Stanley who are told in their morning briefing that the latest statistics available suggest that 92% of black residents in their neighbourhood have open arrest warrants. Later, they go on their separate patrols and each see a black resident. They both form the belief that the black resident has an open arrest warrant. Stella uses this belief to stop the resident and run their name through the database to see if they should be taken in to the precinct. Stanley, too believes that the resident he encounters has an open arrest warrant: however, he chooses not to stop the resident for other reasons. Stanley knows that the justice system is corrupt and that his other police officers have been using warrants and fines on the residents as a way of padding the municipal budget. So, although he believes that this resident has an open arrest warrant, he doesn't act on the belief.[7]

The difference in moral blame between Stella and Stanley suggests that it is not *just* what they believe that is morally problematic. If it were the case that the one belief that some resident has an open arrest warrant were morally problematic, then both Stella and Stanley would be just as bad as one another. But, Stanley seems commendable whereas Stella does not. Thus, it seems that what distinguishes them is how they *act* on the basis of their beliefs. But, if we focus on just the acts, however, we'll fail to capture other ways in which we can wrong one another. In contrast with the case given above, now consider the hermit in the woods who holds racist beliefs.

> *The Racist Hermit.* The racist hermit in the woods will never interact with the disadvantaged person he believes something negative of, he will never interact with or contribute to the institutional structures of racism. He may be a product of these institutional structures, but we are hard pressed to say that he contributes to them given his isolation. But, suppose that he comes to believe that Sanjeev smells like curry. How did he come to form this belief given his isolation from society? Let us just suppose that he discovered

[7] This statistic is loosely based on the Department of Justice's investigation of the Ferguson police department which found that 62% of residents had open arrest warrants, and of those residents, 92% were black. On Aug 25, 2015, a municipal court judge in Ferguson issued an order to withdraw all arrest warrants issued in Ferguson before Dec 31, 2014 due to the Ferguson police and the city's municipal court pattern and practice of discrimination against African-Americans. If you are worried that this case seems contrived, I note that in neighboring predominantly black municipalities, the ratios are even worse. For example, in Country Club Hills there are 33,102 active warrants for a population of 1,381. Similarly, in Wellston there are 15,000 outstanding warrants for a population of 2,460.

> some trash on the ground which happened to be an alumni newsletter from Sanjeev's university that included a picture of him. Upon seeing that picture, the hermit believes of the pictured person—Sanjeev—that he smells of curry. Now, suppose also that Sanjeev happens to have recently made curry so in this instance the hermit's belief is true—Sanjeev *does* smell of curry. Has the hermit wronged Sanjeev?

According to the downstream theorist, the hermit does not and cannot wrong Sanjeev because Sanjeev will never learn of the hermit's belief. However, neither harm nor knowledge are a prerequisite for being wronged. As I will argue further in Chapter 2, the wrong here appears to be *relational*. For example, if Sanjeev's partner were to cheat on him but he never found out, it is a commonly accepted intuition that Sanjeev would still have been wronged. Following Thomas Nagel (1970), wrongs, such as betrayal, needn't be apprehended by the agent that is wronged. To continue pressing this point that we care not only how people act towards us and what they say about us, but also the attitudes they hold of us, consider the following related case.

> *The Security Guard*. Jake is a security guard at a fancy department store. He hates the company he works for, and he couldn't care less if people shoplift and cost the company money. One day, as Jada leaves the department store Jake works at, he believes that Jada shoplifted the purse she's carrying. But, given his contempt for his company, he chooses not to intervene. He acts exactly as he would act if he had believed that that was Jada's purse and that she hadn't stolen it. Has Jada been wronged by what Jake believes of her?

I want to press the intuition that Jada has been wronged by what Jake believes of her. If we put ourselves in Jada's shoes, we don't want people to believe of us that we stole the purse. We care what people believe about us.[8] The problem, then, with downstream accounts of the wrong is that they are at best partial explanations of the wrongs of racist beliefs. It

[8] If you are having difficulty with this intuition, consider the following analogous case from the movie *Pretty Woman*. Vivian Ward, played by Julia Roberts, is an escort who has been hired by Richard Gere's character, Edward Lewis. Edward gives Vivian money for a new wardrobe. Vivian visits a store in Beverly Hills, and the first shot when she enters the store is the reaction from the workers in the store. Although she reminds them she has money to spend in the store, the clerks refuse to believe her and ask her to leave. Although it seems like the wrong committed in this case concerns how the store clerks *acted* towards Vivian, you need only ask almost anyone you know with a darker tint what it feels like in a store where people believe you will shoplift. They may not always follow you around, they may not always act any differently towards you, but that they believe you will steal, or that they believe you can't afford anything in the store, etc. *hurts*. See also Williams (1992, pp. 44-46)'s retelling of her experience in a Benetton store. I discuss this in more detail in Chapter 2.

is right that *one way* people wrong each other in virtue of how what they believe about each other is that how they *treat* each other is a reflection of what they believe about each other. However, the ways in which we can treat each other poorly include not only our words and our deeds, but also what we believe of each other.

Earlier I had summed up the reasons in favour of locating the wrongs of racist beliefs downstream from belief with the following slogan: beliefs don't wrong people, people wrong people. We can still preserve some of the intuitive force of this slogan while nonetheless recognizing that in the cases given above you are wronged by what someone believes of you. We must simply recognize that one of the ways in which people wrong people is by having beliefs about people, not simply what they do or say.

1.3 Locating the Wrongs Upstream From Belief

I have now shown the shortcomings of an account that locates the wrong of racist beliefs solely in the downstream effects of the belief. However, those wanting to resist the thesis that beliefs wrong have the option to locate the wrong elsewhere still, namely upstream. In this section, I turn my attention to accounts that attempt to locate the wrong of racist beliefs *upstream* from belief. Upstream features of a racist belief include the many processes that contribute to the agent forming the belief. On such a view, we could resist the thesis that beliefs can wrong, and instead say that what goes wrong in the case of Spencer is the result of his cognitive predispositions or limitations. Alternatively, perhaps the ways he gathers and evaluates evidence, i.e., his epistemic practices, are flawed. Alternatively, still, the flaws lay in problematic motivations and affective attitudes that are guiding his belief formation. On such an account, beliefs aren't themselves wrong, rather they are but a symptom of a set of states and processes that result in the belief. Further, although this set of states and processes that result in the belief might itself include other beliefs, the problematic beliefs are merely symptoms of a defect in the cognitive system. We can illustrate this distinction between locating the wrong downstream from belief and locating the wrong upstream from belief with the following diagram.[9]

[9]Thanks to Jessie Munton for first illustrating this distinction in a similar way when presenting comments on this chapter at Athena in Action.

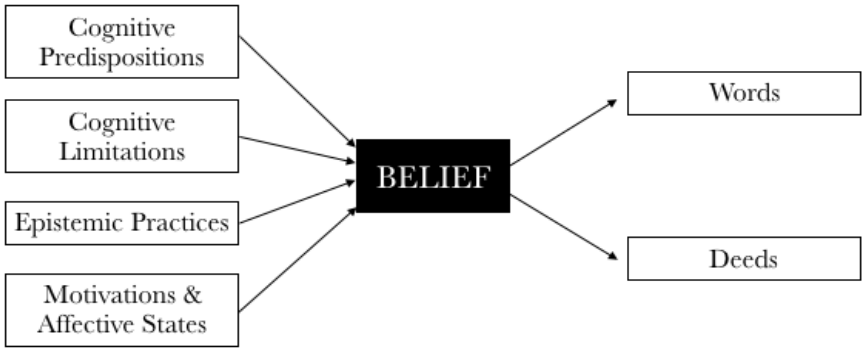

Fig. 1.

We can see an example of this upstream strategy in the case of stereotypes. Stereotypes are helpful in guiding us through the world. As McGarty et al. (2002) note, stereotypes are not only aids to explanation, they are also energy-saving devices or heuristics. Not only are stereotypes devices that we form to help us understand the world, but also they also help us understand the world by saving us time and effort. After all, by treating people as members of a group rather than individuals, we can save energy by ignoring the diverse and detailed information about each. However, because we live in a complex social environment, by taking shortcuts we often adopt biased and erroneous perceptions of the world. That is, rather than being aids to understanding, a common thought is that stereotypes are aids to misunderstanding *more often* than they aid in understanding. So, although stereotypes can assist with explanation, they problematically produce falsehoods and distortions due to upstream features of our cognitive system.[10]

Further, stereotyping not only problematically produces falsehoods and distortions, it does so in a way that is so systematic that it cannot be trusted. Sarah-Jane Leslie (2015), for example, argues that stereotypes are more likely to form about some group if that group is perceived to share an underlying *essence*. For Leslie (2015) these "striking-property generalizations"—e.g., sharks attack bathers, mosquitos carry West Nile—are formed because they treat a group as having an underlying essence that predisposes them to some activity or character trait. Further, our cognitive system is such that we more readily form and endorse these generalizations when they are negative. In sum, the stereotype itself may be unproblematic. There could, perhaps, be creatures who could unproblematically believe stereotypes concerning groups.[11] But given our own cognitive limitations and predispositions, stereotypes wrong. Further, it is not the belief in the stereotype that

[10]This is a simplified discussion of stereotypes and stereotyping. For a more nuanced and careful discussion, see Beeghly (2015) and Antony (2016).

[11]As a note of clarification, the beliefs that are my focus are beliefs about individuals. These beliefs might be formed on the basis of stereotypes. But, I take a belief in a stereotype to be a belief about a group, not a belief about an individual.

wrongs, rather, the wrong lies upstream from the belief.

For another example, perhaps the belief itself does not wrong but given our cognitive predispositions and limitations, we can reasonably expect that the belief will interact with other frame-working beliefs, cognitive biases, etc. in highly systematic and structuring ways. This kind of wrong is tied closely to the the essential functional role of belief. Our beliefs prepare us for more beliefs of the same kind. That is, our beliefs guide us in assessing new evidence as it comes in. This is how priors work. If Spencer has good reason to think that his priors are going to lead him to update on evidence in morally problematic ways, then his failure to appreciate this risk is the wrong he commits. What he believes about black diners in his section is what we could call a *risky amplifying belief*. The belief itself is of neutral value, but once it enters a prejudiced cognitive system, it will have negative effects. Further, the fault lies not with the belief, but rather the prejudiced cognitive system.[12]

You don't have to be explicitly racist or sexist or any other kind of –ist to see this amplifying effect. We live in a world that has been shaped by and continues to be shaped by racist attitudes and racist institutions. So, no matter how well intentioned we are when we add the belief, we may end up with a racist amplified belief set. To see this, consider the cognitive saint—someone who doesn't have any racist background beliefs, assumptions, or biases—and as such can add the belief *this black diner will not tip well* or *Sanjeev smells of curry* without any untoward consequences, and perhaps ought to. The worry is that the mere addition of the belief could cause the entire framework to become biased. The addition of the belief might itself mark the beginning of an implicit association of the sort "black=poor" or "black=bad". As IAT studies have shown us (see Banaji and Greenwald 2013), many people in the US have negative associations of this type towards African Americans. Although, there is still a lot of work to be done to figure out the exact nature and role that these biases play in the cognitive framework, it is not unreasonable to suspect that increased exposure to these sorts of beliefs can create piggybacking associations of the form that I'm concerned about.[13]

In light of these examples, we might want to say that what we believe isn't bad on its own. But, given the way these cognitive biases work they're often invisible to us. So, we have reason to be skeptical about whether we can add the belief unproblematically. So, we should exercise extreme caution (up to and including avoidance) in adding these sorts

[12]Note that a risky amplifying belief could also be a stereotype, but not all risky amplifying beliefs are stereotypes. I owe the term "risky amplifying belief" to a conversation with Renee Bolinger.

[13]See for example Mandelbaum (2016). For some of the controversy surrounding the IAT see Oswald et al. (2013) Oswald et al. (2015), Greenwald and Banaji (2015), and Madva and Brownstein (2016). Thanks to Gabbrielle Johnson for many conversations about implicit bias.

of things to our own belief set. A problem, however, with this account is that it's not clear we can clearly demarcate the wrong *upstream* or whether the wrong is in what we might call the *whole stream*. That is, the wrong lies in the evidence that sparks the belief and how that belief interacts with upstream features of an agent's cognitive framework.

A further problem is that one of the reasons that the thesis that beliefs can wrong seems so unintuitive is because it clashes with the following commonly accepted thesis about control, responsibility, and belief: moral obligation requires a kind of voluntary control that we lack with regard to our beliefs. As I'll show in Section 1.5, although this assumption is often treated as a settled matter of orthodoxy it is up for debate whether we have sufficient control over our beliefs (or whether control is even need) for deontic notions to apply to belief. However, if we take this intuition about control and responsibility seriously, then although it commonly stands in the way of the plausibility of the thesis that beliefs can wrong, here it actually pushes us towards locating the wrong in belief rather than in upstream features that causally impact the belief.

Presumably we have less control over the upstream features than what we lack with regard to our beliefs. As noted earlier, the way in which our cognitive predispositions, limitations, biases, etc. work are often invisible to us. Believing that Jamal will tip less well than James may lead us on a path that makes it easier to form negative beliefs about black individuals we interact with. We have less control over what happens in our cognitive system after we believe than we do before we settle the belief. The upstream theorist would say that the bad thing is whatever cognitive framework allows Spencer to (on the basis of adopting the belief) make a bunch of other terrible decisions. But, the bit that wrongs is the framework, not the other bad beliefs or decisions that result on the basis of it. But if the motivation for locating the wrong in the framework rather than the belief is that beliefs *don't seem* like the kind of things that wrong, the same can be said for an agent's cognitive framework. How does a biased framework, by itself with no beliefs entering as inputs, wrong? This upstream account, thus has only a partial explanation of the wrong committed by Spencer. The wrong is not only in his cognitive limitations and predispositions, but it is also in what he believes.[14]

[14]For more on locating the wrong of racist and prejudiced beliefs in *upstream* features of an agent's cognitive architecture see Begby (2013) and Munton (MS). Begby argues that the reason why prejudices are epistemically insidious concerns how they become interalized as background beliefs that are recalcitrant to empirical counterevidence. Munton argues that the flaw in well-founded statistically-accurate beliefs about certain demographics concerns how we tacitly ascribe the domain of the statistic in question. That is, there are often flaws in our implicit representation of the generality of statistics concerning the rates of violent crime, etc. As noted, I am in agreement that these accounts all get something right, they simply lack the resources to address the character of Spencer, and the peculiarities of the character of Spencer are my primary concern in this chapter and this dissertation as a whole.

1.4 A Restricted Account of When Beliefs Can Wrong

According to a standard way of thinking about the moral wrong of racism, when the wrong is located in belief there are two conditions that the wronging belief must satisfy. The first is the false condition—the belief is false—and the second is the hot condition—the false belief is accompanied by other negative qualities that the belief exhibits, e.g. being held in bad faith, being self-deceptive, being indicative of ill-will on the part of the believer etc. We can already see how such an account will not be able to explain what is morally objectionable about the supposedly rational racist. Although Spencer's belief seems objectionable, it is not false. Spencer's belief appears to be an accurate reflection of the weight of the evidence. Further, Spencer claims to harbour no feelings of ill will.

I grant that in the vast majority of cases people who claim to be like Spencer satisfy either the false condition or the hot condition. The problem, then, with the accounts discussed in this section is not that they are incorrect. Rather, these accounts are incomplete. They foreclose the possibility of a character like the supposedly rational racist. But, as I argued in Section 1.1.2, there is good reason to think that there could be people like Spencer. Thus, my goal in this section is not to argue *against* these accounts. Rather, my goal is to show the limitations of these accounts when they are faced with the particular challenges raised by a character like Spencer.

Let us begin with a prominent account of the wrong of racist beliefs from Kwame Anthony Appiah. Appiah (1990, 1995) argues that racism consists of two elements: *racialism* and *racial prejudice*. Racialism is a proposition about the existence of different races or racial essences. Racial prejudice is a disposition about the moral significance of a race because of some morally relevant racial properties. For example, one might think that some races are more primitive or barbaric than others. The belief rests on a false belief about the existence of different races or racial essences, and it also involves thinking that there are morally relevant racial properties, e.g. being more primitive or barbaric. According to this account, racial prejudice requires accepting a racialist belief, so one cannot have racial prejudice without accepting racialist beliefs.[15] The wrong of racism, then, is that it places unjustified moral significance on race. That is, the racist uses race as a basis for making moral distinctions, e.g., the racist uses race illegitimately as the sole basis for

[15] Although Appiah himself does not think that these racial essences are biological, this account does seem to paint the racist as a biological essentialist. But, that is not necessary. For example, imagine a racist who doesn't believe that the differences are biological, merely cultural. For example, the racist who refuses to hire a Korean dogwalker because of the false belief that Koreans eat dogs. Again, we have a racial proposition *Koreans eat dogs*, which is false, and the disposition to treat their race (or culture) as being of moral significance because it indicates some morally relevant racial property. Alternatively, a clearer example of the cultural racist can be seen in cases involving resentment towards new immigrants from a particular country from the descendants of immigrants from that same country.

treating some people morally. Further, akin to the upstream accounts discussed in Section 1.3, part of the wrong of racism comes from the racist disposition that imbues in the believer a tendency to accept, morally and theoretically, false propositions or beliefs that are hurtful, discriminatory, bigoted, etc. towards one race and not the other. This moral wrong leads to other bad epistemic practices such as accepting racist beliefs in the face of overwhelming evidence that should have led to giving up on those beliefs. In this regard, the wrong of racism is located in the beliefs of the agent, not their actions. Racism, thus, is a cognitive incapacity (see Mills 1997).

Similarly, Lewis Gordon (1995, 2000) argues that racism is a character flaw, but specifically it is a flaw in one's beliefs. Gordon's account, then, is a bridge from the purely cognitive incapacity account of racism—where at its core the wrong of racism lies in a false belief and how that belief interacts with other beliefs—to an account that encompasses the other attitudes of the agent. For Gordon, racism is a character flaw because it is a reflection of one's defective psychology, personality, or character. After all, if it were a mere matter of a false belief, then we wouldn't be able to distinguish the merely ignorant racist who does not deserve moral approbation from the racist that does. Racism must be accompanied by some negative non-cognitive attitudes. This leads us to J. L. A. Garcia (1996, 1997, 1999)'s suggestion that racism is a personal phenomenon of ill-will. This ill-will is often manifested in actions and perhaps helps create and sustain social institutions that contribute to the wrong of racism. The wrong of racism, thus, is this racial disregard. Here, we also hear echoes of Arpaly and Schroeder (2014)'s conception of the wrong of racism: those who have prejudices display some distinctive kind of epistemic irrationality that can be explained by a deficiency of good will. That is, they are ready to believe anything about a particular group as long as it is bad. The racist is guilty of a kind of *hot irrationality*.

To summarize, according to accounts that locate the wrong of racism in what an agent believes, there are two distinct but closely related suggestions. First, the wrong is that the belief is false. Second, the wrong is not a mere false belief, but rather the false belief accompanied by ill-will, hatred, antipathy, etc. I call the first *the false condition*, and the second *the hot condition*. We have already seen that Spencer is not guilty of the false condition. We have been granting that Spencer's belief is true. We have been granting that it *seems* to be rational for Spencer to believe that the black diner in his section will tip less well than the white diner in his section. Perhaps, however, Spencer is guilty of the hot condition. Perhaps his belief is epistemically irrational because it is the result of negative non-cognitive attitudes or ill-will. It is this diagnosis that I turn to considering now.

Following Nomy Arpaly (2004, pp. 102), let us start with the following case.

> Boko Fittleworth (a character in a P.G. Wodehouse novel) overpowers and traps a man whom he spots hiding in his would-be-father-in-law's garden shed at midnight, because he believes this man to be a burglar. In fact, the man is not a burglar but a business tycoon whose presence in the shed is part of a secret, unlikely, and harmless plot in which the future in-law is a willing participant.

In most circumstances, if all that a person gets wrong are the facts, then her false beliefs—i.e. her ignorance—should excuse her action. Although Boko owes the businessman an apology, he is not blameworthy for his action because he has the excuse "But I thought you were a burglar." The racist, however, does not presumably have a similar excuse. Consider, for example, the anti-semite who has the false belief that Jews are members of a worldwide conspiracy set on world domination. Arpaly argues that the anti-semite does not have the analogous excuse "But I thought you were a member of worldwide conspiracy set on world domination." This is an asymmetry in need of explanation. That is, we need to explain why some false beliefs excuse actions, whereas other false beliefs do not.

Arpaly argues that the answer has to do with a difference in the *epistemic rationality* of Boko's belief as opposed to the *epistemic irrationality* exhibited by anyone's belief, in this day and age, that Jews are part of worldwide conspiracy set on world domination. Whereas it is epistemically rational to believe of a man hiding in your shed at night and who refuses to identify himself that he is up to no good, she argues that unless you have just arrived on Earth from another planet with a seriously flawed travel guide, it is difficult to reach the belief that all Jews are involved in a worldwide conspiracy set on world domination. Whereas the first is an *honest mistake*, it is difficult for the latter belief to be an honest mistake. This is because many of the people who hold such beliefs have met Jewish people, know of historical events motivated by such beliefs, and are able to see the unlikely nature of such a conspiracy. So, rather than being an honest mistake, their beliefs seem to be the result of motivated irrationality.[16]

[16]She reiterates this point in Arpaly and Schroeder (2014, pp. 234) when she asks, so, what's the difference between the ordinary racist, i.e. the one who deserves condemnation, and the alien racist who does not?

> One difference that springs to mind is the fact that the Earthly racist holds his belief in, e.g. the supremacy of the majority Chinese ethnicity against plentiful evidence that is readily available to him. With a typical level of intelligence and with the information that is available to rather uneducated people in the first world, he would probably not have developed his racist beliefs if there were not something amiss with him epistemically. The run-of-the-mill racist is epistemically irrational, as are other run-of-the-mill prejudiced people.

This, however, is still not enough. Although some irrational beliefs are morally objectionable, Arpaly recognizes that this does not hold in general. She argues that the person who irrationally believes of their lottery ticket that it will win is not being morally vicious in the same way as the anti-semite.[17] If irrationality in a belief is not sufficient for moral condemnation, what is? Arpaly and Schroeder (2014, pp. 235) argue that the difference is that the bigot's irrationality is "hot" irrationality.

> We agree with Kwame Anthony Appiah that the racist's irrationality is motivated irrationality. His thinking that the Jews are a conspiratorial people is more likely to be caused by his hatred of them, or of people who are ostentatiously "other," or to be caused by his resentment of his low social station, or the like. And the fact that the racist's belief appears to be caused by a desire opens the possibility that, however involuntary, the belief might have something to do with ill will or moral indifference.

According to this account, prejudiced beliefs are morally vicious when they are held as a result of moral indifference. Returning to our challenger, Spencer, he might not think that he is a racist or a bigot because he is just believing in accordance with the evidence, but that is not *why* he believes. If that is what he reports, he is guilty of bad faith or motivated irrationality. He is just ready to believe anything of a particular group as long as it is bad.

This is a strong claim to make about the psychology of such persons. It puts us in a difficult position when arguing with the supposedly rational racist. Firstly, it makes for an uncharitable exchange. Second, when they inevitably deny the description, we're led to a sort of foot-stomping on which neither can convince the other. We would be in a stronger dialectical position if we were to grant that they are not believing in bad faith, or due to motivated irrationality, and still point out how their belief wrongs. This, recall, is the motivation for the working constraint on acceptable answers to the supposedly rational racist, i.e., that we grant Spencer his conception of himself.

Not only would we be in a dialectically advantageous position in response to the supposedly rational racist if we meet this constraint, but also there is another problem with restricting the class of beliefs that can wrong to beliefs that exhibit hot irrationality. An over reliance on the hot condition can obscure other cases of beliefs that wrong. To show this, I demonstrate that there can be an agent with a good will and who has well-intentioned beliefs who nevertheless makes a distinctive kind of mistake when settling a belief, which

[17] Of course, we could add more to the case in which they would be morally vicious.

doesn't seem excused and is morally objectionable. For such a case, consider the following scenario.

> *The Questioning and Unquestioning Scientists.* Sierra and Tango are conducting a study to see if there is a relationship between race and IQ. According to their results, not only is there a 15-point difference between the average IQs of white and black test takers, but IQ is 60% inheritable within the White population. Tango doesn't like the conclusion, but as a scientist he resolves that facts are facts, and he thus unquestionably accepts the conclusion. Sierra similarly does not like this result, but on the other hand is willing to question the result. So, she does some extra research and notices another hypothesis that they did not consider: that the environmental conditions could account (and do a better job accounting) for most of the disparity rather than genetic factors.[18]

Is the badness of Tango's belief due to a deficiency of good will? Arpaly and Schroeder (2014) would say that it is, but I do not find that a convincing gloss on this scenario. It does appear in this case that it is Sierra who is the one with scientific integrity, whereas Tango's acceptance of the result may be due to his prior negative beliefs about IQ in the black population. However, if we alter the scenario slightly, we'll see why that cannot be the whole story.

Let's now set the example in the heyday of scientific racism. That is, scientific experiments are being run in order to justify the racist practices of a society. Now, suppose that the result that Tango* and Sierra* arrive at is the correct result—that there is no relationship between race and IQ. Now, again, Tango* doesn't like the conclusion, but as a scientist he resolves that facts are facts and he accepts the conclusion. Sierra* also does not like the result, but she is willing to question it. So, she does some extra research in order to find a relationship between race and IQ. In this case it is Tango* who exhibits scientific integrity and it is Sierra* that is engaged in motivated reasoning.[19]

I present this pair of cases to illustrate how problematic it is to assume that Tango's acceptance of the results is due to bad faith or motivated irrationality. This is also the point that Spencer is trying to make when he insists that he is not reasoning in bad faith nor

[18] This case tracks the debate between Herrnstein and Murray (1996) and Block (1996).

[19] Further, as I will discuss in Chapter 3, if the amount of evidential support is the crucial factor to determining whether a belief is justified, we cannot distinguish Sierra and Sierra*. The difference between them is that one has an impermissibly low threshold for justification, one that is not matched to the moral stakes of her environment. To explain this, I will argue that we must reject another standard piece of orthodoxy: that sufficiency of evidence supervenes on what your evidence is. As I argue in Chapter 3, the moral stakes can affect whether you have sufficient evidence and whether a belief is justified.

does he have any ill will towards black diners. It is unfair to assume of the supposedly rational racist that he believes what he does because he is ready to believe anything negative about some group, rather than he believes what he does because it is what the evidence suggests. What the case of Tango* shows is why it's wrong to believe of Tango that he only accepts the result due to his prior negative beliefs about IQ in the black population. The supposedly rational racist sees himself as Tango (and in turn as Tango*). He is willing to accept unpopular conclusions if that's what the evidence shows. When it comes to responding to the evidence, Tango and Tango* do better than Sierra and Sierra*. In contrast, our moral intuitions suggest that Sierra is better than Tango, but also that Tango* is better than Sierra*.

Further, it is now well-documented that although it is tempting to think that the racist suffers from ill will or a deficiency of good will, racial injustice can survive and even thrive in the absence of such negative non-cognitive attitudes. Racism can come in *cold* varieties as well as *hot* varieties. So, the hot condition is not necessary for identifying the moral wrong of racism.[20] This is why it is not only a dialectical advantage if we can provide an account that doesn't undermine Spencer's conception of himself, this really ought be a constraint on any acceptable answer to what Spencer does wrong when he believes what he believes.

1.5 Another Hurdle: A Worry About Control

My aim in this chapter has been to present a case—the case of the supposedly rational racist—that pushes us to seriously rethink why we reject the naïve thesis that beliefs can wrong. I've argued that this case forces us to reconsider why we are so resistant to the idea that I could wrong you by believing something of you that I seem to be justified in believing. The challenge posed by the supposedly rational racist is a challenge concerning our epistemic practices in a non-ideal world. A common refrain throughout this chapter has been the following: the world is an unjust place and there may be many morally objectionable beliefs that it justifies. As a result, the evidence might be stacked in favor of racist beliefs. The world we inhabit is a racist one, so it is no surprise then that some of our beliefs are racist as well.

If our best accounts of the way racist beliefs can wrong require the believer to always act on the belief before we call the belief wrong or require the believer to be distorting the truth, misrepresenting reality, organizing facts in a misleading way, or engaging in

[20]See for example Valian (2005)'s account of the gender disparity in philosophy, which she calls a "cold" social-cognitive account given that "it is purely cognitive rather than emotional or motivational" (pp. 1).

fallacious reasoning, then these accounts will fail to meet the challenge posed by the supposedly rational racist. Spencer appears to be responding to the evidence in a way that seems epistemically rational. This is why I suggest that we stop looking elsewhere for what wrong Spencer commits. The question I am interested in is what happens when we take seriously the thought that Spencer wrongs the black diner in his section in virtue of what he believes of the diner. This is the question that will motivate the rest of this dissertation.

There is, however, another hurdle to clear to establish the plausibility of the thesis that beliefs can wrong. This hurdle is the assumption that moral obligation requires a kind of voluntary control that we lack with regard to our beliefs. As we have seen, this background assumption is the motivating force behind the attempts to locate the wrong committed by Spencer either in the downstream effects of the belief or somewhere causally upstream from the belief itself. My argument against this assumption is going to be short and indirect. As I shall show, there is more intuitive plausibility to the thought that the supposedly rational racist wrongs given what he believes than there is to the thought that voluntary control is required to wrong.

The problem of control can be presented as follows. Although we may talk in everyday discourse about what people should believe, deontic notions such as *ought* and *should* and moral notions such as *praise* and *blame* do not apply in the domain of epistemology. The argument goes as follows.

1. For deontological concepts such as obligation or duty to apply in the domain of epistemology, agents need to have voluntary control over their beliefs.

2. It is not the case that we have sufficient voluntary control over our beliefs for deontological concepts such as obligation or duty to apply in the domain of epistemology

3. Therefore, deontological concepts such as obligation or duty do not apply in the domain of epistemology.

To briefly canvas various arguments against the idea that we either lack control or that we even need control for these deontic notions to apply, Nishi Shah (2002) rejects (2) on the grounds that we exert control over our beliefs in the following way: through our appreciation of the evidence we are agents with regard to our beliefs and are capable of regulating our beliefs. Matthias Steup (2000) similarly argues that our beliefs are responsive to epistemic reasons and in this way we exercise control over our beliefs. Further, Brian Weatherson (2008) has persuasively argued that this argument relies on too narrow a conception of voluntary, and Amy Flowerree (2016) presents the following reductio to

establish that we are agents with respect to our beliefs: if we are not agents with respect to our beliefs, then we are not agents with respect to our intentions. Others have also argued that (1) is false, i.e. for deontological concepts to apply, control is not needed. For example, Pamela Hieronymi (2006, 2008) argues that it is not troubling to say that we can no more intend at will than believe at will. She further rejects that voluntary control is required for deontic notions to apply. Connor McHugh (2012) similarly argues that voluntariness is not a central condition. In general, these strategies have taken the following form: they show that for deontological concepts to apply in the domain of epistemology, either we have sufficient control or there is no control needed for deontological concepts to apply.[21]

What we are left with is the following choice. Should we give Spencer a pass on the basis of a theoretical technicality that is both (a) not as intuitively obvious as it is commonly regarded, and (b) does a disservice to the common folk-intuition that something has gone wrong in the case of Spencer? Our choice is between saying that Spencer is not wronging at all or that Spencer wrongs through what he believes. The main consideration preventing us from saying the latter is this theoretical normative posit that voluntary control is always required to wrong. However, this intuition is more esoteric and has less claim to be either the intuition of the folk or share widespread acceptance amongst philosophers than the intuition that I have been defending: that Spencer does something wrong.

To conclude, I hope to have convinced the reader that we should take seriously the thought that people wrong people in virtue of what they believe about each other, and not just in virtue of what they do. To further show *how* beliefs can wrong, is the next task, and to answer that we turn to the next chapter.

[21] In Basu and Schroeder (2018), we also consider another way of resisting the siren's call that is the problem of control. We reject the following underlying assumption that motivates the problem of control: the thesis that our beliefs are at the mercy of the evidence. I will also explore this is in more detail in Chapter 4.

CHAPTER II

What We Epistemically Owe To Each Other

At the end of the previous chapter we ended on the following note: Spencer does *something* wrong and that whatever that wrong is, that wrong is properly located in what he believes of the diner in his section. The task of this chapter is explain what it is that Spencer does that is wrong. That is, to explain how, why, and when beliefs—beliefs of the sort Spencer holds of the diner in his section, i.e., beliefs about another person–wrong. Notice that I have not said that the task is to explain *what* beliefs wrong. As I shall show, the wronging is not a matter of the *content* of the belief, but rather concerns a morally problematic way of *relating* to others through what we believe of them. As a result, the title of this dissertation is misleading. You might expect that a dissertation titled "Beliefs That Wrong" would tell you *which* beliefs wrong. Instead, as we'll see over the course of this chapter, a better name would've been "Doxastic Morality". My goal in this chapter is to show how moral demands that are constitutive of the relationships we stand in with one another give rise to moral demands not only on our words and deeds, but also our beliefs.

In addition, another goal of this chapter is to show that it is not only racist beliefs that can wrong in this way, but that many of the beliefs that we form about one another can wrong. The particular beliefs held by the supposedly rational racist that were the focus of the previous chapter are merely an instantiation of a wider class of beliefs that wrong. As I shall show, when we relate to others and form beliefs about them in the same way we relate to objects and causal phenomenon and form beliefs about tables and the weather, we fail to give others what they are owed: recognizing them as agents rather than objects. In general, we wish to be seen as individuals, not as objects to be predicted and managed. This failure to relate as we ought is a morally criticizable failure that stems from a kind of moral indifference, a kind of moral disregard, for which we can be held accountable. Already you may object that this thesis seems to make an implausibly wide range of things morally problematic. So, another task of this chapter is to show that it's not an implausibly wide

range of things that are morally problematic, rather, it's perfectly plausible that the wide range of things this chapter discusses are in fact morally problematic. Thus, the three tasks of this chapter are (i) give an explanation of why Spencer's beliefs wrongs, (ii) show that Spencer's moral failing is an instantiation of a more general way in which we can wrong others through what we believe, and (iii) address the worry that this proposal makes a wide range of things morally problematic.

Finally, in way of set-up, I should note that this chapter will be primarily case driven. In Section 2.1 I will present a number of cases that suggest that there is an overlooked aspect—the cognitive or epistemic aspect—of the moral demand we place on one another to be treated well. These cases—*Mistaken Identity, Wounded By Belief, The Racist Hermit, The Security Guard,* and *Sherlock Holmes*—all suggest that there is a moral demand, a demand to be related to in a way that is characteristic of being an agent rather than an object, and this moral demand has an *epistemic* dimension. We fail to relate as we ought when we try to make predictions of other agents as though they are objects whose acts are determined by causal laws, e.g., beliefs about someone's status as a diner or a waiter on the basis of their skin color. In Section 2.2, I turn to presenting a picture of morality that can explain how moral considerations matter to the question of how we settle belief. I will argue that in virtue of our relationships *and* in virtue of moral features of our environment, e.g., structural injustice, we fail in what we epistemically owe to each other when we observe and form beliefs about other people in the same way we observe and form beliefs about planets. In Section 2.3, I note that although there is this general moral concern we all have in virtue of being agents—the wish to be seen as individuals, not as objects to be predicted, anymore than managed—for some groups of people, in particular marginalized groups, this epistemic-moral demand is particularly strong.

Finally, in Section 2.4 I attempt to provide some guidelines for recognizing the conditions under which there is a greater risk that we will fall short of our epistemic obligations to one another. That is, I attempt to articulate some general principles for when beliefs wrong. As we'll see, however, this is no easy task. The diversity of cases that push us seriously considering the thesis that beliefs can wrong also makes it difficult to articulate any principle (or set of principles) that can provide necessary and sufficient conditions for identifying when beliefs wrong. To avoid wronging an agent must not only be aware of certain moral facts, but also many contingent social, structural, and historical facts as well that determine when those moral considerations are applicable. In short, the agent must exhibit a moral-epistemic virtue that is akin to *being woke*. The account of moral encroachment to be defended in Chapter 3 will be an attempt to further flesh out what this moral-epistemic virtue looks like. But first, let us begin by considering some cases.

2.1 Examining Cases

2.1.1 Mistaken Identity

> *Mistaken Identity.* The conference has ended, and the organizers have had the forethought to book a number of tables at a nearby restaurant so that conversation can continue over dinner. You're having a good time at dinner and, after a few drinks, you get up to use the restroom. As you return to your table, one of the diners, Jim, reaches out to grab your arm and says, "I asked for a refill fifteen minutes ago." For a moment you're confused, then it dawns on both of you what mistake has been made. Most philosophers don't look like you. With regard to melanin levels, you share more in common with the wait staff than your fellow diners. Given your skin colour, the likelihood that you are a member of the staff rather than a fellow diner was high enough to seemingly make it rational for Jim to assume that you were a waiter, not a fellow diner. The belief that Jim had—and in turn his actions—might amount to a social faux pas, but, given that the belief was well-supported by the evidence, the belief—and in turn his act—was reasonable. He's not a bad guy; he just made an honest mistake. Of course, in the moment you don't reason through all of that. In the moment, you quickly laugh it off, eager to return to your table and dinner conversation.

Later, you notice that you can't stop thinking about that interaction, and you wonder why such a small thing is still bothering you. It's not just the act that's bothering you: something about your colleague's *belief* that you were a waiter is keeping you up at night. So, you lay in your bed and you run the incident over and over again in your head. You think to yourself, well, some epistemic mistake was made—the belief was false—the belief was well-supported by the evidence. Presumably, given that the belief was well-supported by the evidence, the belief was also rationally permissible. If all of the of dinner guests, except you, are white whereas all the staff members share your skin complexion, then based on race alone, it seems that anyone in the position of Jim—needing to flag down a waiter for their drink order—should similarly believe that you are a staff member. That is the doxastic attitude that is seemingly supported by the evidence.

But, you also think to yourself, isn't it also the case that mistaking a person of color for a member of the staff is one of the paradigmatic examples of racism? In *Slumdog Millionaire*, Jamal is mistaken for a tour guide while he himself waits for a tour of the Taj Mahal. In *Miss Saigon*, Kim is mistaken for a maid when she enters a hotel room. Barack Obama has often been asked to get coffee even when dressed in a tuxedo at a black-tie

dinner.[1] The examples are numerous. Given how strong the association is between *non-white* and *staff*, it's not surprising that when a young Whoopi Goldberg saw Lt. Uhura—the communications officer on Star Trek played by Nichelle Nichols—she ran around the house screaming "there's a black lady on television and she ain't no maid!"[2] In light of this, let us reconsider the opening case.

Returning to our dinner companion, we can imagine Jim offering the following apology: "My evidence supported that you were a waiter, not a conference participant, so I had the attitude I epistemically ought to have had towards you being a staff member as supported by my evidence." Further, he might add, "When I believed you were a staff member I did not actively feel any ill will towards you, nor exhibit indifferent disregard."[3] This excuse, however, does not assuage the feeling of resentment one bit. You are still mad at him. And it is not just his action—flagging you down as a waiter— that you are mad at, it's that he *believed* that you were a waiter. Imagine if things had gone differently, and Jim still had the belief but didn't act on that belief. After all, a common response to these cases is that Jim should believe, but also that he should not act on his belief. Even in such a case, there seems to be something upsetting about the way in which he so readily believed that you are a waiter. Consider an analogous case. Suppose, for example, someone saying to you that they believe the stereotype that women are bad drivers, but that they would never act on that stereotype. That does not make the fact that they believe that women are bad drivers any less insulting. The belief remains insulting even when the communicative act is removed.

It would be all too easy to dismiss these feelings as irrational. I will use this case, and others, to argue that it is not irrational to be upset. Jim has done something wrong, and that one has these feelings of hurt, anger, indignation, and resentment when these events happen should be taken seriously. I will argue that we should take these feelings seriously, and my argument follows the Strawsonian thought that our reactive attitudes such as resentment are often good indicators that some moral demand hasn't been met and the Lordeian thought that our feelings are a source of evidence and information.[4]

[1] As Obama notes, "there's no black male [his] age, who's a professional, who hasn't come out of a restaurant and is waiting for their car and somebody didn't hand them their car keys." See: http://people.com/celebrity/the-obamas-how-we-deal-with-our-own-racist-experiences/

[2] https://www.yahoo.com/tv/blogs/tv-news/exclusive-video-whoopi-goldberg-star-trek-213710247.html

[3] Regular folk don't talk in this way, but remember, we're imagining this case taking place at a post-conference dinner.

[4] There is already a lot of ground to cover in this chapter, so to restrict it's scope, let me note that I will just be assuming that we can rely on our feelings in this way. One might object that anger is not the appropriate response because anger is not virtuous. It might be reasonable or excused to feel anger, after all, it can be morally exhausting to be magnanimous or forgiving or to laugh off the fifth time you've been pulled over

2.1.2 Wounded By Belief

> *Wounded By Belief.* Suppose that Mark has an alcohol problem and has been sober for eight months. Tonight there's a departmental colloquium for a visiting speaker, and throughout the reception, he withstands the temptation to have a drink. But, when he gets home his partner, Maria, smells the wine that the speaker spilled on his sleeve, and Mark can tell from the way Maria looks at him that she thinks he's fallen off the wagon. Although the evidence suggests that Mark has fallen off the wagon, would it be unreasonable for Mark to seek an apology for what Maria believes of him? He is, after all, wounded by what his partner believes of him.[5]

I contend that cases like *Wounded By Belief* are common place and such cases point us to a folk practice of holding people accountable for their beliefs, for thinking that it is the beliefs themselves that wrong, and for being upset with others for what they believe about us. Consider, for example, how common it is to say the following, "You shouldn't have believed that of me."

One may try to respond by objecting that Mark is not wronged until he has clarified that he hasn't had a drink. Perhaps Mark has only been wronged if *after* he has clarified that he hasn't had a drink Maria continues to believe that he has had a drink. Further, the reason why Mark deserves an apology is because of Maria's act of discounting his testimony, not merely her having the belief. I grant that this could be a partial explanation of the wrong. However, we can imagine that even if Maria does not discount Mark's testimony, a wrong still occurs. Mark should be able to demand that his partner think better of him and not immediately settle on the belief that he must've had a drink—even when the evidence strongly suggests that he has. He would be *further* hurt if she continues to believe that he's had a drink after he explains that he hasn't, but there is still the initial hurt before any exchange occurs.

As I will argue in Section 2.2, Mark has been wronged *because* his partner has taken an agent-neutral stance towards forming beliefs about him. Further, that some people do weaponize and abuse this move in relationships—demanding that their partner believe better of them, disregard or discount evidence, etc.— provides additional support

by police or mistaken for a staff member. Nussbaum (2016) puts forward such an argument. But, as her critics note, anger plays an important role as a response to injustice in the world and provides importance information and evidence about injustice (see Bell 2009, Srinivasan 2016, 2017, Bommarito 2017, and of course Lorde 1984). For the purposes of this chapter I will just follow the latter's assumptions about the role of emotions in guiding inquiry into injustice.

[5] See Basu and Schroeder (2018).

for thinking that they are only able to effectively to do so if it is in fact a practice that we expect within a relationship.[6]

One might further object to this case that what we are reacting to in this case is not what Maria believes of Mark, rather, we are reacting to a violation of trust. Through her belief Maria reveals that she doesn't trust Mark to not have had a drink. What Mark then seeks an apology for, what he wants to be corrected, is that trust that Maria has violated. That is where the wrong is located; the wrong is not in the belief, the wrong is in the violation of trust.[7] I grant that the belief may be an indication of a lack of trust, and it is that lack of trust that Mark, and in turn if the case were about us, we would be wounded by. But, it does not follow from that that we are not also wounded by the belief. Here is an analogy to press this point. A slap in the face may be indicative of someone's ill will towards you. One thing we desire in our relationships with others is that they not feel that ill will towards us. Nonetheless, we are wounded by the slap, and can rightly seek an apology for the slap. Similarly, the belief may be indicative of a lack of trust, and what we want is that there be trust in our relationship. Nonetheless, we are wounded by the belief, and can rightly seek an apology for the belief.

I grant that one may remain unconvinced. But to see how deep this folk practice goes, let us consider more cases.

2.1.3 The Racist Hermit and The Security Guard

Previously, in Chapter 1, I presented the following two cases: *The Racist Hermit* and *The Security Guard*. Whereas in both *Mistaken Identity* and *Wounded By Belief* the belief in question seems *epistemically justified* but false, in these two cases the beliefs are not epistemically justified but in one case the belief is true and the other the belief is false. Nonetheless, in both cases the beliefs seem to wrong. In Section 2.1.4 we will consider a case that parallels that of Spencer, the supposedly rational racist: a case in which the belief seems epistemically justified and the belief is true yet the belief seems to wrong. The purpose of these cases is to illustrate the diversity of circumstances in which we seem to wrong others given what we believe. Sometimes we wrong when the belief is false, sometimes when it's true, sometimes when we're justified and sometimes when we're not. The task after these cases are presented is to then determine what is common to all the cases. But first, let us consider *The Racist Hermit* and *The Security Guard*.

[6]Thanks to Elle Benjamin for pressing me on this point. See also Stroud (2006)'s related work on the moral demand on friends that they believe better of their friends. I will discuss Stroud's work in more detail in section Section 2.2.2.

[7]Thanks to Laura Gillespie for pressing me on this point.

> *The Racist Hermit.* The racist hermit in the woods will never interact with the disadvantaged person he believes something negative of, he will never interact with or contribute to the institutional structures of racism. He may be a product of these institutional structures, but we are hard pressed to say that he contributes to them given his isolation. But, suppose that he comes to believe that Sanjeev smells like curry. How did he come to form this belief given his isolation from society? Let us just suppose that he discovered some trash on the ground which happened to be an alumni newsletter from Sanjeev's university that included a picture of him. Upon seeing that picture, the hermit believes of the pictured person—Sanjeev—that he smells of curry. Now, suppose also that Sanjeev happens to have recently made curry so in this instance the hermit's belief is true—Sanjeev *does* smell of curry. Has the hermit wronged Sanjeev?

In this case the racist hermit's belief is true, and it in no way affects Sanjeev. Nonetheless, I suggest that we have an intuition here that the hermit is still doing something wrong. If you think that only actions or words said to a person can wrong, then in this case you'll have to say that the hermit does not (nor cannot) wrong Sanjeev because Sanjeev will never learn of the hermit's belief. But, as I noted earlier, when it comes to *relational harms* neither harm nor knowledge of the harm or wrong are required to be wronged. For example, if Sanjeev's partner were to cheat on him but he never found out, it is a commonly accepted intuition that Sanjeev would still have been wronged. For another case to continue to warm the reader up to the idea that we really care what people believe of us and can be wronged by what they believe of us, consider the following case.

> *The Security Guard.* Jake is a security guard at a fancy department store. He hates the company he works for, and he couldn't care less if people shoplift and cost the company money. One day, as Jada leaves the department store Jake works at, he believes that Jada shoplifted the purse she's carrying. But, given his contempt for his company, he chooses not to intervene. He acts exactly as he would act if he had believed that that was Jada's purse and that she hadn't stolen it. Has Jada been wronged by what Jake believes of her?

I want to press the intuition that Jada has been wronged by what Jake believes of her. If we put ourselves in Jada's shoes, we don't want people to believe of us that we stole the purse. We care what people believe about us. The ways in which we can treat each other poorly include not only our words and our deeds, but also what we believe of each other.

Finally, before one dies of exhaustion from cases, one may be unconvinced by *Wounded By Belief* because you might think that *if* Mark had been drinking, he wouldn't have been wronged by Maria's belief. Hence, what is doing the work in getting the intuition that Mark has been wronged is that Maria *falsely* believes that Mark has been drinking. Note, however, that in *The Racist Hermit* the belief is true, nonetheless there is the intuition that Sanjeev has been wronged. To further press this point, consider a case that is analgous to the motivating case of the dissertation, i.e., Spencer, the supposedly rational racist. Recall that Spencer's belief about the diner in his section seems to be both epistemically justified (and we can assume for this purpose) a true belief. To see an analogous case, let us consider Sherlock Holmes.

2.1.4 Why Everyone Hates Sherlock Holmes

Within the fiction people are constantly annoyed with and get mad at Sherlock Holmes for various reasons. Of particular interest to this chapter, Sherlock Holmes is an excellent example of someone who engages in this kind of morally objectionable statistical reasoning that we find in the previous cases. Further, he engages in this kind of reasoning with regard to every person he meets. Further, without fail, the people he encounters find the beliefs he forms about them to be insulting. Sometimes, for one in Victorian England, it's because it is a negative belief, e.g., that someone is sexually promiscuous on the basis of mannerisms that increase the likelihood that they are sexually promiscuous. But other times, people also respond negatively to his belief about what they ate for breakfast. What remains in common is the manner in which Sherlock Holmes forms beliefs about people. He observes everyone as objects to be studied, predicted, and managed. I suggest that the negative way in which people respond to what Sherlock Holmes believes about them, whether it concerns their culpability in a crime or what shoe they put on first when getting dressed in the morning, can be explained by feeling wronged when people look at us the way that Sherlock Holmes looks at us.

As noted, within the fiction people are constantly annoyed with and get mad at Sherlock Holmes for various reasons. Of particular interest for this chapter and this dissertation, often the annoyance stems from some yet-unarticulated moral demand with regard to forming beliefs about others that Sherlock Holmes consistently fails to meet. Sherlock Holmes, as a character, presents us with a really suggestive case in support of the claim there's a way of viewing people—observing people in the same manner we observe natural phenomena like planets, the weather, etc.—that wrongs the people viewed in such a way. Consider, for example, for the following scene from the 1985 movie, *The Young Sherlock Holmes*.

SH: Wait–let me. Your name is James Watson. You're from the North, your father's a doctor, you spend much time writing, and you're fond of custard tarts. Am I correct?

JW: My name isn't James, it's John.

[...]

SH: Very well, so your name is John. How did I do on the others?

JW: You were correct. On every count. How is it done? Is it some sort of magic trick?

SH: No magic, Watson. Pure and simple deduction. The name-tag on your mattress reads "J Watson". I selected the most common name with "J". "John" was my second choice. Your shoes aren't made in the city. I've seen them before when visiting the north of England. Your left middle finger has a callus, the trademark of a writer. You were carrying a medical book not available to the general public, only to physicians. Since you can't have been to medical school, it was given to you by an older person, someone who is concerned for your health: Your father, the doctor.

JW: And the custard tarts?

SH: Simple. There's a stain of yellow custard used in making tarts on your lapel, and your shape convinced me you've eaten many before.

JW: There's no need to be rude.

Now Sherlock, of course, engages in this kind of reasoning with regard to every person he meets. And, without fail, the people he encounters find the beliefs he forms about them to be insulting. This negative way in which people respond to what Sherlock Holmes believes about them, whether it concerns their culpability in a crime or what shoe they put on first, illustrates how we feel wronged when people look at us the way that Sherlock Holmes looks at us, how we feel wronged by the way in which Sherlock Holmes forms beliefs about us.[8]

[8] A worry one might have here is that this wrong that Sherlock commits, this failure to relate to others as he ought, is something that is impossible for him to correct. That is, folks who are neurotypical may be able to relate to others in morally unproblematic ways that are impossible for Sherlock Holmes. Perhaps Sherlock cannot but relate to others as though they are objects or causal phenomena, and as a result, a moral theory that concludes that Sherlock is constantly wronging others seems perverse, it compounds Sherlock's plight. This objection is part of a larger problem in moral philosophy concerning neurodiversity, disability, and the demands of morality. For the time being I must set this issue aside, but I want to thank Regina Rini for pressing me to think more about this issue.

2.1.5 Reflections

What we have now seen is a wide array of cases. As noted earlier, in some we wrong when the belief is false, in some when the belief is true, in some the agent seems epistemically justified and in others the agent doesn't. See, for example, the following table.

	Ep. Justified?	True?	Negative Bel?	Moral Wrong?
Mistaken Identity	Yes*	No	Yes	Yes
Wounded By Belief	Yes*	No	Yes	Yes
Racist Hermit	No	Yes	Yes	Yes
Security Guard	No	No	Yes	Yes
Sherlock Holmes	Yes*	Yes	No	Yes
Spencer	Yes*	Yes	Yes	Yes

Table 1.[9]

What is common to all of these cases is this nagging intuition that we can be wronged by what others believe of us. Further, the the beliefs we are wronged by needn't be negatively charged beliefs, i.e., something that most people would not want believed of them. The beliefs could be mere "observation[s] of trifles."[10]

Over the course of this chapter I will be arguing that a commonality in these cases is that when someone looks at us not as a person but as though we are an object that is determined by causal laws, as something whose behaviour is to be predicted, that is to step back from seeing us as a person. Naturally, that is upsetting. You might now question whether we can move from the mere fact that these scenarios are upsetting to the conclusion that someone has been wronged. As I will argue, I think we can, and my goal is to show how this could be true. My goal, now, moving forward into Section 2.2 is to show that it is not idiosyncratic to think that moral considerations matter for belief. After all, our relationships are partly constituted by how we form beliefs about one another. Although when we think of ethics, we think of moral theories that concern maximizing the good, or willing into being universal laws, and the such, there is a tradition at the heart of moral philosophy that doesn't limit the scope of what we care about to just how people act towards us and say of us, but also what others think of us.

[9] I have put a '*' by the "Yes"'s because, as I'll explain in Chapter 3, the beliefs aren't really epistemically justified. They could potentially be epistemically justified, but the wrong in these cases is that the threshold for justification does not correctly match the moral stakes of the environment.

[10] See Doyle, *The Boscombe Valley Mystery*.

2.2 On How We Relate to Each Other

That the way we ought relate to each other is importantly different from the way we relate to objects is not a new idea. It can be found in the Kantian idea that we need to recognize each other as legislators in a kingdom of ends acting under the idea of freedom, the Sartrean idea that due to our consciousness we transcend the facticity that is our situation and are thus *more* than our situation, the Darwallian idea that by standing in relevant authority relations with each other we can give each other second-personal reasons or the Strawsonian idea that we just happen to be the kind of beings that care deeply about the attitudes that we hold about each other. Through weaving together a story about stances and standpoints that draws on these various accounts and philosophical traditions concerning how we ought relate to one another, I aim to demonstrate that there is a cognitive dimension to how we ought relate to one another.

As I stated earlier, my goal is to show that this idea is not idiosyncratic, i.e., that all these various accounts and philosophical traditions are intelligibly gesturing at the same moral ideal with regard to how we relate to one another. Further, that whatever the story about *why* we owe what we owe to each other, there's a cognitive element to what we owe. So, over the course of this section I develop an account of what I call *the moral standpoint* that draws on these various accounts. Importantly, the account need not commit itself to additional claims concerning the nature of free will, the human being as a being-for-itself, etc. As I'll show, the moral standpoint encompasses not only what we owe each other in word and deed, but also what we to owe each other in thought.

2.2.1 The Involved and The Participant Stance

Let us start with a Kantian idea; as Rae Langton (1992) explicates it there is an important way in which we relate to other people that is characteristically different from how we relate to objects. This idea stems from the following picture of the world and our place in it. We, as human beings, find ourselves in a world that consists of things—tables, chairs, corn, cotton—and we try our best to understand how this world works, e.g., why plants grow when watered, why dogs only give birth to other dogs and never cats, etc. To understand the world we turn to science as a way of discovering these patterns in nature. But, we not only try to understand how things work, we think also about how to use them, i.e., how they can be used as a resource, as a means for human ends. Crucially, according to this Kantian picture, each of these things can be priced, and each is essentially

replaceable.[11] But, the world consists not only of things; in it there are also people. It is then, "in our dealings with people, [that] we have a different way of going on, though it is hard to capture just what that is" (Langton 1992, pp. 486). It is in an effort *to capture just what that is* that Langton turns to Strawson to find the following idea:

> We 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, pp. 486)

What is it, then, to be involved? P. F. Strawson (1962) argues that insofar as we stand in relationships with one another, we are susceptible to a wide range of responses from resentment to love to shame, etc. This is characteristic of *the participant stance*. Of these responses, a narrow class—closely associated with responsibility and blame—are often referred to as the 'reactive attitudes'. These reactive attitudes include resentment, indignation, and guilt. These reactive attitudes are characteristic of *being involved* as they are reactions to the attitudes and intentions people have towards us, and they concern a connection between the attitudes and moral demands.[12] Further, a significant part of being so involved concerns according a certain importance to others' attitudes and intentions towards ourselves and being cognizant of those demands with regard to our treatment of others. Strawson argues that we ought take an involved stance towards others and recognize that, just as our own self-understanding relies on the attitudes that others take towards us, so too do theirs. This is what we owe to each other.

Some theorists (such as Wallace 1994, 2014) attempt to construe this idea of being involved in a relationship quite narrowly to include only the paradigmatic cases such as spouses, friends, lovers, family, etc. But, as we can see in the following quote, Strawson has a much broader notion of relationships:

> We should think of the many different kinds of relationship which we can have with other people—as sharers of a common interest; as members of the same family; as colleagues; as friends; as lovers; as chance parties to an enormous range of transactions and encounters. Then we should think, in each of these connections in turn, and in others, of the kind of importance we attach

[11] Although I am sympathetic to the general Kantian picture, there are many things to dislike about the Kantian picture. For example, it is not the case that all things in the world in virtue of not being people are thereby essentially replaceable. For example, the environment, natural wonders and monuments, etc.

[12] These reactive attitudes "are essentially reactions to the quality of others' wills toward us, as manifested in their behaviour: to their good or ill will or indifference or lack of concern" (Strawson 1962, pp. 194).

> to these attitudes and intentions towards us of those who stand in these relationships to us, and of the kinds of reactive attitudes and feelings to which we ourselves are prone. (Strawson 1962, pp. 187)

We can also see a similar sentiment expressed when Tim Scanlon (1998, pp. 168) notes that morality

> ...requires us to be moved by (indeed to give priority to) the thought of our relation to a large number of people, most of whom we will never have any contact with at all. This may seem bizarre. But if the alternative is to say that people count for nothing if I will never come in contact with them, then surely this is bizarre as well.

The *participant stance* consists in a particular way of caring about or being invested in the attitudes of another person. Further, this is simply what we owe to others in virtue of standing in relationships with them.[13] Further still, these relationships we stand in with one another create the possibility of wrongs that wouldn't exist otherwise.[14]

Moreover, our feelings are indicators that such moral demands haven't been met. Consider, for example, the anger you feel when you call out to a person to hold the elevator door for you and they ignore your calls, the disappointment you feel when you tell a friend a secret and they share it with someone else, or the irritation you feel when someone cuts you off in traffic. We can easily recognize that you have been wronged in some way in each of these cases. As I retell the stories, I start to feel the disappointment, the anger, the indignation, etc. And I am sure that you start to feel the same way. When moral demands or these expectations, such as trust, fair regard, respect, etc. are not met, a sign that they are not met is these reactive attitudes.[15] So, in the spirit of Strawson (and Langton's Kant), I suggest that we take these feelings as indicators that some moral demand has not been met and that we have been wronged. Given that these attitudes are present in the above cases, we can infer that some demand hasn't been met. The question now is what moral demand hasn't been met.[16]

[13] I owe much of my thinking on this topic to Steve Bero and am grateful to him for his discussion of Strawson (1962) in Bero (MS).

[14] This is easiest to see in the case of close relationships, such as the demand for doxastic partiality in friendship. I will address this point in more detail in Section 2.2.

[15] Even if you are the kind of person that strives to not have these negative emotions, they are often an almost automatic response to having been wronged. A drink is spilled on you and immediately you feel anger, though it may subside for various reasons, e.g. you realize it was a child. More on this in Section 2.3.

[16] Marusic and White (2018) develop a similar account to what I will be calling *the moral standpoint* by (i) extending some Strawsonian thoughts in moral philosophy to the realm of epistemology, and (ii) some Burgean apparatus. They argue that we wrong in belief the same way we wrong in action. We wrong, in both cases, when we fall short of justified normative expectations, for example, when we fail to be genuinely

Consider again the person that doesn't hold the elevator, the person that cuts you off, the person that shouts "Gross!" when you sneeze. Their attitudes and their actions convey an indifference and lack of moral care with regard to how they affect. Now, looking back to the opening cases, you may object that the examples that have been used to illustrate the Kantian idea of an involved stance or the Strawsonian idea of the participant stance all concern action. The motivating cases of this chapter, however, concerns thought, in particular, what is believed of you. Perhaps, for example, in *Mistaken Identity* when the diner pulled on your sleeve mistaking you for a waiter, they were showing this criticizable kind of indifference in their action, but not in their belief. We find this same resistance, this same attempt to limit what we owe each other to exclude epistemic demands, in Stephen Darwall (2006). So, I turn now to addressing this Darwallian worry and arguing for an extension of moral care to the realm of thought.

2.2.2 The Cognitive Demands of The Moral Standpoint

Darwall (2006)—in a manner similar to both Kant and Strawson—argues for what he calls *the second-personal standpoint*, i.e., the perspective you and I take up when we make and acknowledge claims on one another's conduct and will. In speech, we can see this perspective in performatives such as demanding, apologizing, etc. In thought, we see this perspective in Strawsonian reactive attitudes like resentment and guilt. The reasons that stem from this standpoint—namely, second-personal reasons—derive from our relations with one another. Further, he argues that our concept of moral obligation is *second-personal* as it concerns a commitment to the idea that those to whom we are morally accountable have the authority to demand that we do what they ask of us. In addition, when we make demands on others we do not always do so as individuals, we do so as representatives of the moral community.

Darwall distinguishes this second-personal stance from the third-person perspective, i.e., where we regard each other "objectively" or "agent-neutrally". Interestingly, for our purposes, Darwall explicitly denies there is a cognitive element to the second-personal stance. Although he argues that the second-person standpoint is the perspective we take up whenever we make valid claims or demands on someone else (whether explicitly in speech or implicitly in thought), he also argues that this epistemic authority is third-personal as opposed to second-personal. Epistemic authority, unlike the relevant author-

cooperative. Belief, they argue (as I will too), is part of the participant stance. hey differ, however, by arguing that instead of adopting evidential policies (as Paul and Morton 2018 argue) or adjusting our thresholds (as I will argue in Chapter 3, and as Basu and Schroeder 2018 and Schroeder 2018b argue), what is required of us within our relationships and our shared activities of reasoning is an epistemic permission or a default entitlement to believe others.

ity for second-personal reasons,

> ...depends fundamentally on a person's relations to facts and evidence as they are anyway, not on her relations to other rational cognizers. Even in cases of testimony where we take someone's word for something, this second-personal authority can be defeated by deficiencies of epistemic authority of the ordinary third-personal kind. If we have reason to distrust her beliefs or judgment, we also have reason to reject her second-personal epistemic claims. (Darwall 2006, pp. 12)

Darwall (2006, pp. 56-7) further goes on to say that,

> [o]f course, what reasons people have to believe things about the world depend in many ways on where they stand in relation to it. But ultimately their reasons must be grounded in something that is independent of their stance, namely, what is the case believer-neutrally. Our beliefs are simply the world (including our place in it) as seen (committedly) from our perspective; what we should believe depends ultimately on the world as it actually is.

There is something right to Darwall's claim here, but there is also something that he misses. What he misses regards how the relationships we stand in with each other change the stance we ought take towards the reasons that ground our beliefs. But first, let me start with what Darwall gets right.

Ordinarily, when it comes to settling the question of whether to believe *p*, all that you should take into consideration is whether *p* is true, and correspondingly, reasons that are likely to get us onto the truth of *p*. In this way, reasons for belief are, as Darwall suggests, grounded in something that is independent of our stance. The role of our beliefs is to accurately capture the world as it actually is, not how we would like it to be. However, even if we accept that we ought believe in accordance with our evidence, we can still ask the question of when the evidence is sufficient to justify belief. We recognize that in some cases we need more, and stronger, evidence than in other cases. For example, we require more evidence when passing judgement in a criminal case than when settling a playground dispute. Although Darwall does note that what reasons we have to believe things depends on the ways and where we stand in relation to it: the reasons we have to believe things depend also on the stance we take towards others. Our standards for justification are not stance-independent. Crucially, in order to meet the moral demands in our relationships, these standards must be stance-dependent.

I suggest that what Darwall misses is that our relationships change how we should respond to evidence. Our reasons for belief with regard to other people must not be

grounded in something independent of the stance we ought take towards others. This point is initially hard to see because it runs counter to our usual intuitions about belief and our belief-forming practices. So, let me illustrate the point with some cases that speak to the intuition behind this point, i.e., the point that we can fail to meet the demands of our relationships precisely by taking this believer-neutral stance.

Consider, for example, the moral requirement that we believe better of our friends. Sarah Stroud (2006) argues that we have special responsibilities toward our friends which we don't have toward strangers that extend beyond our actions towards our friends versus strangers. Being a good friend also requires a kind of epistemic partiality. What we ought believe of our friends crucially cannot be something that is independent of our stance towards them. That is, we should not be neutral with regard to how we respond to evidence about our friends. For example, suppose you receive testimony that Arjun cut someone off in traffic. Ordinarily, that testimony would be sufficient to justify your belief that Arjun cut someone off in traffic. If Arjun is a close friend, however, your threshold for when the evidence is sufficient to justify believing poorly of your friend is much higher. So, although testimony would ordinarily be sufficient to justify the belief, because you stand in this relationship with Arjun you require more evidence before you will be justified in believing poorly of your friend. So, you have a reason for belief that is independent of your stance, but whether that is sufficient to believe is something that is stance-dependent.

Similarly, Paul and Morton (2018) ask us to to consider a person who adopts the goal of finishing a doctoral dissertation. As I find myself in the same position, let's just consider the question of whether I will succeed in finishing my doctoral dissertation. What should you believe about whether I'll succeed? The seemingly obvious answer is whatever the total accessible evidence supports. The *unobvious* answer, they argue, is that it depends on your relationship to me. Normally, when beliefs are influenced by personal relationships, those beliefs are our paradigmatic cases of epistemic irrationality. After all, reasons arising from our relationships are irrelevant to the truth of those beliefs. Perhaps we have reasons arising from our relationships to *act as though* we're more confident in our loved one's ability to succeed, but we shouldn't *believe* that they are more likely to succeed. Paul and Morton argue that although there are not practical or ethical reasons for belief, there are practical or ethical influences on the standards by which we reason about what to believe.

To establish this claim, they note that evidential policies govern the way we adjust our evidential thresholds in different contexts. However, there is no *uniquely* best evidential policy to have. There are multiple evidential policies that are rationally permissible for a given thinker to have from the point of view of purely epistemic considerations. Fur-

ther, if that's the case, then practical and ethical considerations can and should play a role in deciding between epistemically permissible policies. Further, these evidential policies can also be shaped by our relationships with others. Within the context of significant relationships, Paul and Morton argue that the default perspective we should take towards one another is *the rational perspective*. This perspective is contrasted with *the diagnostic stance* in which you view another person as an object whose behavior can be explained and predicted like a mechanism. The rational perspective, on the other hand, requires viewing others as rational beings to be engaged with on rational terms. We see that this distinction tracks the earlier Strawsonian distinction between the objective and the participant stance, and I will argue for a *moral standpoint* akin to Paul and Morton's *rational perspective*. To briefly summarize then, when we doubt our loved ones, we can wrong them if that doubt is the product of an overly low evidential threshold for adopting the diagnostic perspective. In short, what we believe of others is *not* stance-independent.

Finally, Berislav Marusic (2015) also notes that when it comes to beliefs about what we will do in the future, the stance we take to such beliefs is importantly different from the stance we take towards other beliefs. When it comes to what you will do in the future, you should believe against the evidence because whether we succeed is up to you. Further, those who are partial to us such as our friends, lovers, our spouse, etc. would be wronging us if they didn't also believe against the evidence. In virtue of being close to you, they should exhibit a similar kind of doxastic partiality to what we discussed above. As Marusic (2015, pp. 2) notes,

> [w]e cannot be scientific observers of ourselves without distorting our agency. In particular, when something is up to us, we distort our view of what we will do if we seek to predict what we will do—even if our predictions are based on excellent evidence. Similarly, we cannot be scientific observers of others without distorting our relationship to them. In particular, if we are close to them, we distort our relationship if we seek to predict what they will do—even if our predictions are based on excellent evidence. Agency and partiality require a view of ourselves and others that is incompatible with a fully naturalistic view [...].[17]

What we see here is epistemic authority having a second-personal nature. The reasons we have to believe can derive from our relations with one another.

[17] It is worth noting here that Marusic (2015)'s view is also heavily influenced by Kant, Sartre, and Strawson. His positive account introduces Sartre to develop what he calls a Sartrean response—approaching the question of what to do as a practical question, not as a theoretical question—as another way of developing the Kantian dictum that we act under the idea of freedom, as well as Strawson's claim that we take a participant point of view toward ourselves.

Before continuing, let me also briefly note the following. What Marusic (2015) says here seems overly strong. Sometimes part of what is required in close relationships is to predict things about our partners. In virtue of being in a close relationship we are better predictors of our friends, partner, colleagues, etc. Consider the following case. You've been talking with your partner on the phone and they seem curt with you. Do you wrong your partner, do you distort your relationship with them, if you suspect they've been having a bad day so on the way home you stop and buy them something that shows you care?[18] Sometimes the way we show our thoughtfulness is by being good predictors of the people we care about. So, to simply say that the way in which beliefs can wrong is due to a kind of moral indifference that stems of a lack of care would not get us the full picture. Sometimes showing the amount of care that is demanded by the relationships requires treating one another as objects to be managed and predicted. We can accommodate this, however, by simply noting that not all objectification is bad. Sometimes it is licensed, excused, or even invited.[19]

However, if these moral demands emerge from the relationships we stand in with others, we not only have close personal relationships, we also plausibly have relationships with strangers in virtue of being part of the same moral community. Recall the following Scanlon (1998, pp. 168) quote from earlier.

> [Morality] requires us to be moved by (indeed to give priority to) the thought of our relation to a large number of people, most of whom we will never have any contact with at all. This may seem bizarre. But if the alternative is to say that people count for nothing if I will never come in contact with them, then surely this is bizarre as well.

Now consider the following complicating case. Suppose that Maria is riding the bus and a complete stranger gets on the bus smelling of alcohol. Ought she believe that they've been drinking, or should she withhold judgment? I suggest that although the nature of Maria's relationship with the stranger on the bus is different from the nature of her relationship with Mark, there is still something owed in this case. It shouldn't be surprising that we also owe strangers something like moral care when forming beliefs about them. After all, just as we are invested in what they think about us, they are also invested in what we think

[18] Perhaps it's flowers, chocolate, a DVD of a terrible movie, etc. For me, nothing makes me feel better on a bad day than Mexican coke and a bag of XXtra Flamin' Hot Crunchy Cheetos.

[19] Perhaps the Kantian will be helpful again here. When you use your partner's stomach as a pillow, when you manage their emotions when they're hungry, etc. you are not *merely* relating to them as an object. As for the Kantian I can still respect my taxi driver as an end-in-themselves although I am using them as a means of transport between point A and B, perhaps too in these cases when we make predictions of others we can do so in a way that still respects their agency.

about them. The difference, however, between Wounded By Belief and the stranger on the bus is not that in one case Maria wrongs and the other she doesn't, rather it's a point about a difference in degree of the wrong. Maria owes *more* to Mark than to the stranger on the bus. We owe more to some people in virtue of our relationships with them.[20]

To briefly recap before continuing, what I have done so far is provide support for the claim that our reasons for belief, in particular reasons for believing things of others, stem from and are constrained by moral demands that emerge from the relationships we stand in with one another. With regard to believing things of other people, we can wrong them through what we believe if our beliefs are not responsive viewing them as agents as opposed to objects to be predicted or managed. This suggests that a key part of the *wrong* that is done concerns a kind of objectification and alienation. I turn now to Jean-Paul Sartre (1943) to expand upon this idea further.

The basic Sartrean idea is that the human condition is both being-in-itself and being-for-itself. We have both facticity—what we are by our nature, the objective facts about us—and transcendence—an ability to surpass what we are by our nature, to act against it, our consciousness. With regard to the relationships we stand in—being-with-others—we encounter others through the "look". We become aware of ourselves when we are confronted by the gaze of the other. That is, it is when we are watched that we are aware of ourselves (as an object), and this gaze is objectifying. Others look at us and define us as a thing, as a being-in-itself. When seen this way we are objectified. The other robs us of our freedom and our understanding of ourselves as a being-for-itself. This gaze, Sartre argues, is alienating. We seek in our relationships to be seen as we wish to be seen; hence this objectifying gaze of the other unsettles our sense of self.[21]

Combining this with the previous accounts, we can see why it hurts to be seen as a waiter even though there is nothing wrong with being a waiter. Your fellow dinner patron—a person whom you consider an equal, a colleague, someone with whom you

[20]Further, as I shall show in Section 2.3, we also owe more to some people rather than others in virtue of moral features of our environment such as structural injustice. The nature of our relationships, the moral environment, whether the belief is negative, whether the belief involves taking the third-person perspective or the objective stance, etc. are all conditions that play a role in determining whether a belief wrongs and whether we've fulfilled our epistemic obligations to one another.

[21]This distinction also plays an important role in de Beauvoir (1949)'s explanation of women being seen as the other, as a sex object whose being is to be interpreted in relation to men's desires and choices, rather than being seen as a free consciousness. Women's actions are seen as causally determined and women are seen as governed by their facticity by their biology and their hormones. Part of the struggle, then, of existing as a woman is that you are forced to see yourself as both an unfree object and as free. This is a kind of double-consciousness that we see articulated in Du Bois (1903, pp. 2). I will expand on this point in Section 2.3 where I explain why although we in general owe everyone moral care with regard to forming beliefs about them, we also owe more moral care with regard to beliefs about members of marginalized groups.

stand in a special relationship—fails to see *you*. Look again at the apology that's offered: "My evidence supported that you were a waiter, not a conference participant. So, I had the attitude I epistemically ought to have had towards you being a waiter as supported by my evidence. This belief wasn't motivated by any feeling of ill-will toward you, I was just believing in accordance with the evidence." We see in the apology that he takes this theoretical perspective. He is responding as though he is a scientific observer explaining casual phenomenon in the world. But, that is precisely what he should be apologizing for. He should be apologizing for failing to take the participant attitude. For failing to see why you would be invested in the attitudes he holds of you. Similarly, Maria in *Wounded by Belief* fails to recognize that when forming beliefs about Mark in virtue of their relationship she owes him more than an ordinary person on the street (though of course, some is still owed there, just less).

This, I suggest, can be captured under the demands of *the moral standpoint*. We 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. The moral care we must exercise in our beliefs, then, is a reflection of the fact that we should not be as quick to settle our beliefs about other people as we are with our beliefs about natural phenomena. The relationship we stand in with others is importantly different from the one we stand in with objects. This may be due to recognizing each other's humanity, recognizing that we we act under the idea of freedom, recognizing that we are not a being-in-itself but also that we each are a being-for-itself, etc. I suggest that it can also just simply be that we are the kind of beings for whom it matters what people think of us. We 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. There need not be a deeper story than that to ground what we owe to each other.

This is the thought that lies behind the wrong that's committed to Jamal in *Slumdog Millionaire* when he's mistaken for a tour guide, to Kim in *Miss Saigon* when she's mistaken for a maid, to Barack Obama when he is mistaken for a waiter or a valet. If the world is structured such that the evidence suggests you are more likely to be a staff member on the basis of race so that is what people believe of you, we can now see why it mattered so much to Whoopi Goldberg that she could finally see a black woman on television who was playing an officer on a spaceship, not a maid. To repeat the slogan of this chapter, you should not observe another person, form beliefs about another person, in the same way we observe and form beliefs about planets. This is the cognitive dimension of what we owe each other. This is the cognitive demand of the moral standpoint.

2.2.3 The Objective, The Theoretical, The Naturalizing Stance

There is, however, another objection that can be raised here. We are not always required to take the participant stance. How, after all, is sociological work possible if we must always take the participant stance towards others? Sometimes, the perspective of the scientific observer is the view we should take towards others. So, perhaps some of the opening cases are similar, perhaps those are also cases in which we are excused from taking the participant stance.

To understand the objective stance, let us turn back to Strawson. Strawson (1962, pp. 190) notes that to adopt the objective attitude, in contrast to the participant attitude, is to see another person as

> an object of social policy; as a subject for what, in a wide range of sense, might be called treatment; as something certainly to be taken account, perhaps precautionary account, of; to be managed or handled or cured or trained; perhaps simply to be avoided.

We adopt this objective attitude when the agent is excluded from ordinary adult human relationships, for example, if they are a child. It can also be available to us as a resource when aiding in policy or to use out of intellectual curiousity. Although it is often inappropriate to take the objective attitude, sometimes we move between the participant and the objective attitude. For example, the therapist ought to take the objective stance towards their patient, my doctor ought treat me as an object to be managed with regard to risk factors for disease and ailments, a social scientist ought treat her subjects as objects to be studied, etc. So, if the opening case is sufficiently similar to cases in which we ought adopt the objective attitude, then we can disregard the hurt feelings described in the case. For example, when my doctor informs me of my greater chance of heart disease given my racial background, I might be upset at the news, but that is not a criticism of my doctor's judgment. Similarly, if Jim believes that I am a waiter given my racial background, although I might be upset about being mistaken for a waiter, perhaps it is not a criticism of his judgement.

We should note, however, that there are nonetheless still moral constraints on how my doctor is allowed to treat me, how they convey their beliefs about me, and even what they believe of me. Suppose a more serious case than merely increased risk of heart disease. Suppose my doctor informs me I have cancer. It would be inappropriate if they laughed at me while telling me that, or if they announced the news without the seriousness it deserves. Further, although it is appropriate for my doctor to believe that I have a greater chance of heart disease because of my race, it would be inappropriate for them to believe

that I was a bad tipper on the basis of my race—even if the weight of the statistical evidence were the same in both cases.

Consider, for example, a case that Keisha Ray (2017) describes. She is suffering from high blood pressure, so the doctor takes into account her racial background and prescribes a drug that is shown to be effective for black patients. Although it is appropriate to take into account her racial background for that purpose, note the inappropriateness of the doctor adding the following, "Plus it's cheap, so it's good for poor black people." On the basis of the same statistical racial information, one belief was fine—the belief about which drug to prescribe—but the other belief was not—the belief about Ray's financial status. As Ray further notes,

> [e]xperiences like this have made me hyperaware of my blackness when I enter a doctor's office. Often, patients complain of feeling like just another number to their doctors, but when I go to the doctor I feel like just another black person. My doctors don't see me. They don't see that I am a black person who is also a woman, a professional philosopher, a friend, a runner, a candy aficionado and so forth. They just see a black person.

Recall that this is the wrong that I articulated in Section 2.2. We don't want to be seen just as a waiter, just as a black person, just as a shade of brown.

For now, I will just note that taking the objective attitude doesn't free you entirely from moral demands. Here's a general principle that doctors nonetheless follow: serious matters deserve more moral care. Further, if my doctor is licensed to form beliefs about me as a scientific observer might, that is due to the relationship we stand in with one another. The doctor-patient relationship is different from the relationship I stand in with others. For example, whereas it is appropriate for my doctor to believe that I'm at a higher-risk for heart disease given my racial background, it would be inappropriate for a person on the street to believe that. Similarly, in the context of a doctor's office it appropriate to be asked to fill out a form listing your past medical history and your current ailments. It would, however, be inappropriate if an employer asked you to do so.[22] To repeat, occupying the objective attitude does not free one from all moral demands.[23]

[22] Consider the episode "Health Care" from Season 1 of *The Office* in which Dwight Schrute is tasked with picking a new health care plan for the employees. Initially he chooses an incredibly cheap plan with little coverage. In response to backlash from the employees, he then distributes forms that ask all the employees to list the ailments and illnesses they have that they wish to be covered.

[23] As a clarificatory point, I should note that part of the reason why the objective attitude does not completely free us from the moral demands that seem more appropriate for the participant stance is due to the fact that we are still participants in these relationships and ought recognize each other as other agents even when occupying this objective attitude. It may be possible that in some cases, such as the doctor-patient

There is another question here that remains unanswered. I had said that if my doctor is licensed to form beliefs about me as a scientific observer might, that is due to the relationship we stand in with one another. However, with regard to doctor-patient relationship, is it only the *relationship* that creates the privilege of taking the objective stance? That is, is it that the relationship changes things such that it's consistent with the engaged stance to view the evidence in this detached theoretical way? Alternatively, is it the subject matter—your body and how it operates, as opposed to social factors concerning poverty, crime rates, etc.—that changes what is owed? That is, perhaps forming beliefs about others on the basis of some statistical evidence, e.g., medical diagnoses that attend to purely biological factors (for now we are excluding the role that social factors have on people's health), is morally fine, whereas other forms of statistical evidence, e.g., social facts, are problematic. In the abstract, I think there is something to this idea. Perhaps there are morally neutral reference classes we can invoke and form beliefs about individuals on the basis of such reference classes. However, our world is messy. As I noted in my earlier parentheses, a doctor who attends only to biological markers and ignores a patient's social setting, family and generational history, etc. will have an incomplete profile of their patient's health. This is why I think the approach of trying to find morally unproblematic reference classes that we can reason with to be a non-starter.

There are still two more complications that arise out of what I've said so far. Previously I noted that whereas it is appropriate for my doctor to believe that I'm at a higher-risk for heart disease given my racial background, it would be inappropriate for a person on the street to believe that. However, there are at least two additional cases in which taking the objective stance may be what's morally required. First, sometimes we must recognize objective facts about others. For example, I am a visible minority in more than one regard in philosophy. As a result, you should recognize that I've likely had certain experiences in the profession that others who do not share my demographics have not had. If you were to disregard that, you would fail to understand me. Consider, for example, the following poem by Pat Parker (1978), "For the white person who wants to know how to be my friend".

> The first thing you do is to forget that I'm Black.
>
> Second, you must never forget that I'm Black.

To understand each other, we must not forget that human beings aren't just free to define

relationship, you could take a completely objective attitude towards another person and thus not feel compelled by moral demands. But, to take such a stance would be to stop seeing the other person as human. To stop relating to another person and recognizing them as a human agent is to wrong them.

themselves as they wish.[24] Sartrean *bad faith*, after all, involves either acting as though you are totally free to define yourself despite how others see and interpret you or acting as though you are the role you've decided to play. Human beings are complicated, and how we ought relate to each other is similarly complicated. What we must hedge here is what Schroeder (2018a, pp. 5) calls the *Kantian Gambit*.

> *If and when* you interpret someone's behavior in causal terms, this explanation goes, you take the objective stance toward them, and *if and when* you take the objective stance toward someone, you thereby preclude the participant stance

As Schroeder goes on to note, sometimes treating someone as a person is compatible with treating them as a thing, and sometimes treating someone as the thing that they are is how we successfully engage with and understand each other. As I will argue in Section 2.4 and take up in more detail in Chapter 3, perhaps given these complications there are neither necessary nor sufficient conditions we can state for when beliefs wrong. Instead, the best we can do is develop our epistemic character such that it responds correctly to the morally relevant features of our environment in a virtuous manner. As I'll be noting later, we already have reason to think there exists a moral-epistemic virtue of this kind: *being woke*.

For now, let us turn to one final complication. That is, another case in which perhaps what we morally ought do is take the objective stance not only towards others, but also towards ourselves. And I warn the reader now that this discussion will concern both rape and torture. Strawson himself notes that sometimes we can take occupy the objective stance as a respite from the strain of human engagement. In addition, we can imagine cases in which, for reasons of self-protection, perhaps we ought not relate to the best interpretation of another's behaviour. For example, during traumatic events, depersonalizing the other (or even oneself) might be a matter of self-preservation. For a less-morally-laden example, imagine that going to the dentist is an anxiety-inducing prospect. In such a case, it can be helpful to not think of the dentist as another person or to even consider yourself as a person. Rather, you are an object being operated on. In cases where a much deeper moral wrong is being committed to you or against you, e.g., cases of torture or rape, relating to the other as required by the participant stance would compound the wrong.

As David Sussman (2005) has argued, part of what makes torture a distinctive moral wrong concerns the interpersonal relationship it enacts. Effective torture pits the victim

[24] Thanks to Wendy Salkin, Serene Khader, and Aarthy Vaidyanathan for pressing me to say more on this.

against herself. The victim, Sussman (2005, pp. 25) argues, "is forced into a position where she must try to anticipate and understand every little mood and quirk of her torturers. Despite herself, she finds herself trying to grasp her torturer's interests, anticipate his demands, and present herself in a way that might evoke pity or satisfaction from him."

Similarly, consider Offred and The Commander in Margaret Atwood's novel The Handmaid's Tale. Every month there is a *ceremony*, that is, a monthly rape of Offred by The Commander. The ceremony is bearable only if Offred pretends that it is not happening to her, if she is able to be detached from both The Commander and his wife. But soon The Commander begins inviting Offred to play scrabble with him because the impersonalness of "the ceremony" bothers him too and he wishes to befriend her. He tries to earn her trust, he gives her gifts, he allows her to read in his presence (an act that is forbidden). These acts of taking the participant stance, of becoming engaged with and relating to others as persons now makes the ceremony an even more troubling case of wronging. As Offred remarks.

> When the night for the Ceremony came round again, two or three weeks later, I found that things were changed. There was an awkwardness now that there hadn't been before. Before, I'd treated it as a job, an unpleasant job to be gone through as fast as possible so it could be over with. Steel yourself, my mother used to say, before examinations I didn't want to take or swims in cold water. I never thought much at the time about what the phrase meant, but it had something to do with metal, with armor, and that's what I would do, I would steel myself. I would pretend not to be present, not in the flesh.
>
> This state of absence, of existing apart from the body, had been true of the Commander too, I knew now. Probably he thought about other things the whole time he was with me; with us, for of course Serena Joy was there on those evenings also. He might have been thinking about what he did during the day, or about playing golf, or about what he'd had for dinner. The sexual act, although he performed it in a perfunctory way, must have been largely unconscious, for him, like scratching himself.
>
> But that night, the first since the beginning of whatever this new arrangement was between us—I had no name for it—I felt shy of him. I felt, for one thing, that he was actually looking at me, and I didn't like it. The lights were on, as usual, since Serena Joy always avoided anything that would have created an aura of romance or eroticism, however slight: overhead lights, harsh despite the canopy. It was like being on an operating table, in the full glare; like

> being on a stage. I was conscious that my legs were hairy, in the straggly way of legs that have once been shaved but have grown back; I was conscious of my armpits too, although of course he couldn't see them. I felt uncouth. This act of copulation, fertilization perhaps, which should have been no more to me than a bee is to a flower, had become for me indecorous, an embarrassing breach of propriety, which it hadn't been before.
>
> He was no longer a thing to me. That was the problem. I realized it that night, and the realization has stayed with me. It complicates. (pp. 160-161)

These cases can be taken in two ways. First, they could provide more evidence for the moral demand that we relate to others as persons, that we occupy the participant stance, the rationalist perspective, the moral standpoint, etc. That is, without this requirement that we relate to others in a particular way, we wouldn't be able to explain this way in which torture is a distinctive kind of moral wrong or how when The Commander and Offred begin to recognize each other as people, no longer as objects, things become complicated. But, it also suggests that there may be times when it is morally required not to take such a stance towards others. After all, if a moral theory requires that I relate to my torturers and believe the best of them, that is a strike against the moral theory.

In short, when we are morally required to take the objective stance is a complicated matter. It is not clear that we'll be able to articulate a general principle for when we can recuse ourselves from the demands of engagement that the participant stance and the moral standpoint require. Human life is complicated, there are cases in which we ought not relate to others as persons, but instead see them and engage with them as things. In general, however, the point still stands that our default attitude towards others should be that of the moral standpoint. The cases discussed in this section should certainly not be regarded as the norm.

2.3 When We Epistemically Owe More to Some Than To Others

I have focussed on how our relationships can change what we epistemically owe. And as I suggested at the beginning of this chapter, it is not only our relationships that create moral demands on what we ought believe of others, so too do moral features of our environment such as structural injustice.

To articulate why there is this demand and the differing degrees it comes in, let me turn to a very general observation about the human condition that lays behind Strawson's thoughts about the participant stance.[25] We each conceive of ourselves under a variety of

[25]Here I am indebted again to Bero (MS).

more or less stable conditions, from marginal ones such as being born on a Friday, to quite central ones such as being a philosopher, being a spouse, etc. The more central self-descriptions are important to our sense of self-worth, to our self-understanding, and constitute our sense of identity. They are the way in which we understand who we are and where we stand in the world. When these central self-descriptions are ignored in favour of what we are expected to be on the basis of our race, on the basis of our gender, on the basis of our sexual orientation, etc. we see ourselves being reduced to this Sartrean *being-in-itself*. This is a particularly severe wrong in the case of non-dominantly situated groups.

For example, imagine mistaking a white man at a Beyoncé concert for a staff member rather than a concert-goer. A mistake has been made; a wrong has been done. You failed to relate to him as he sees himself, and instead formed a belief on the basis of well-founded statistical evidence concerning the racial demographics of concert-goers and staff at a Beyoncé concert.[26] I suggest that although there is a wrong in this case, it is much less severe than the wrong to marginalized groups. The difference depends on the way in which members of marginalized groups are *dispositionally vulnerable* with respect to their self-descriptions while members of dominantly-situated groups are not. Members of marginialized groups are more dependent on this Sartrean idea of "the look" or "the gaze" than members of dominantly-situated groups are.

As Desirée Melton (2009) notes, to be dispositionally vulnerable involves an awareness of self dependence on others for understanding and respect. Yes, we are all invested in how other people see us, but in the case of racial and gender oppression, of homophobia, of being disabled, you are more dependent than if you are not a member of a marginalized group. Your sense of self is more *fragile*. We owe everyone moral care, but we owe some people more than others due to the way they are more dependent on others for their self-descriptions. To help illustrate this point, consider W. E˙ B. Du Bois's concept of double-consciousness. Du Bois (1903, pp. 5) notes the following.

> It is a peculiar sensation, this double-consciousness, this sense of always looking at one's self through the eyes of others, of measuring one's soul by the tape of a world that looks on in amused contempt and pity. One ever feels his two-ness,—an American, a Negro; two souls, two thoughts, two unreconciled strivings; two warring ideals in one dark body, whose dogged strength alone keeps it from being torn asunder.

When you are not a member of a dominantly-situated group, you must see yourself through

[26] Thanks to Gabbrielle Johnson for suggesting this case.

two self-conceptions. First you see yourself how the dominantly-situated group sees you; second you see yourself how you wish to see yourself. Importantly, one of these conceptions is maintained structurally, i.e., it matters not whether any individual believes that you are lazy, there exists this structurally maintained conception of members of your racial group. An inner conflict arises when those two perspectives inevitably clash. Recall the quote from Ray (2017) earlier. Ray sees herself as a professional philosopher, a friend, a candy aficionado, etc. All that her doctor sees is a black patient. Further, all her doctor sees is a patient, who in virtue of being black, is likely also poor. This is the difference between the white man at the Beyoncé concert and the black man at a fancy restaurant. The first person might be mistaken for a staff member, but given the structure of our society, he does not suffer from this feeling of double-consciousness. As Simone de de Beauvoir (1949) notes, the white male perspective is considered the universal perspective. For the second person—for Barack Obama, for Whoopi Goldberg—when they are mistaken for staff members, that deeply challenges their central self-descriptions from which they draw their sense of self-worth.

To conclude in a manner similar to Kant, Strawson, Sartre, Marusic, et al., I suggest that in many of the circumstances we find ourselves in we ought not be scientific observers towards other people. When we are, we distort our relationship to them. Fulfilling what we owe to others involves an epistemic dimension. This moral demand is the demand to be related to in a way that is characteristic of being an agent, rather than an object. And for some groups of people, in particular non-dominantly situated folks, this moral demand is particularly strong. We ought not relate to other agents in the same way we relate to determined causal phenomena. We all wish to be seen as individuals, not as objects to be predicted and managed. This is how morality can make demands on belief. This is how beliefs wrong. This is the epistemic dimension of what we owe to each other.

2.4 Recognizing Our Epistemic Obligations

My goal has been to make plausible the claim that fulfilling what we owe to others involves not only being sensitive to how we act towards each other and what we say of each other, but also what we believe of each other. Along with that theoretical task of clearing space for this kind of view I favour because it makes sense of our intuitions about the cases in Section 2.1.1-2.1.5, I wish to conclude by leaving the reader with some practical advice with regard to recognizing when we're most at risk of failing to meet our epistemic obligations to one another.

In many of the circumstances we find ourselves in we ought not be scientific observers

towards other people. I have noted, however, that although there are circumstances in which we ought, the nature of the relationship and the moral demands of the environment determines when that's the case. To determine when this is the case, I suggest the following guidelines.

> In general, a belief *p* wrongs a subject *S* when some or more of the following conditions are met,
>
> (i) the belief is a threat to their moral standing,
>
> (ii) the belief is a negatively charged belief,
>
> (iii) the belief coheres with an existing social stereotype that is being invoked, and/or
>
> (iv) the belief involves viewing someone as an object, as something to be predicted and determined, unless otherwise licensed by the relationship.
>
> ...but not always

With regard to (i), I've suggested that part of why we owe marginalized groups more turns on the way that what we believe of them can be a threat to their moral standing insofar as we understand their sense of self as a key part of their moral standing (as I think we should). With regard to (ii), however we want to cash out that condition, it should be clear that seemingly positive beliefs can also wrong, e.g., supposedly "positive stereotypes" such as African-Americans excel at athletic activities and Asian-Americans excel academically can wrong the individuals believed about because their agency is reduced to statistical generalizations concerning their group membership.[27] With regard to (iii), in many of these cases where we fail in what we epistemically owe there exists a social stereotype that is being invoked, and the belief falls in line with that stereotype. Perhaps it is a greater threat to one's standing when that happens. But, there need not necessarily exist a social stereotype in the background for the belief to wrong. And finally with regard to (iv), although it is often the case that viewing someone as an object, as something to be predicted and determined, wrongs them, I've also shown that that is not always the case.

One way, then, of determining the degree to which a belief wrongs is by seeing how many of the conditions I have outlined hold with regard to the belief. But, no condition

[27] Perhaps we should instead consider what Schroeder (2018b) calls *diminishment*. Schroeder argues that beliefs wrong only when they *diminish* you, i.e., when they *bring you down*. To bring someone down, Schroeder argues, is to interpret them in a way that makes their contribution out to be less. To be diminished, then, is to be seen as *less than*, to not been seen as a full agent. Consider again the positive stereotype that South Asians excel academically. I may have benefitted from this stereotype, but there is also a sense in which when people believe that I will excel academically *because of* my ethnicity, my successes are *diminished*. My successes are no longer a reflection of *me*.

by itself settles the degree to which the belief wrongs, or even in some cases whether the belief wrongs. But, they're still helpful guidelines. Racist beliefs, for example, satisfy all of these conditions. As a result, that simply suggests that racist beliefs—the kinds of beliefs that cohere with an existing social stereotype, the kinds of beliefs that view another person as an object, the kinds of beliefs that are negative, the kinds of beliefs that are a threat to a person's moral standing—are the worst kinds of beliefs. Such beliefs are the most morally bad of the beliefs you can hold of another person. That is in fact what we think of racist beliefs.

Note, however, that calling out someone for their white privilege might also satisfy all four of these conditions.[28] In her discussion of white fragility, Robin DiAngelo (2011, pp. 54) presents the following case.

> [She has] just presented a definition of racism that includes the acknowledgment that whites hold social and institutional power over people of color. A white man is pounding his fist on the table. His face is red and he is furious. As he pounds he yells, "White people have been discriminated against for 25 years! A white person canÕt get a job anymore!"

We see that he is reacting in anger. Acknowledging one's white privilege is a threat to one's moral standing. If being good, kind, compassionate, egalitarian, etc. is central to your conception of your self, then being told that you are a racist and that you are complicit in racist structure is a threat to that understanding of yourself. But if anger is a guide to moral wronging, in what way has he been wronged here? Are all threats to one's moral standing, one's sense of self, problematic? Similarly, consider the following case. A black kid is walking the streets of Chicago when he sees a white male police officer. The kid then forms the belief that the police officer is a bad cop who will (or has) probably engaged in police brutality against other black residents in Chicago.[29] Here the cop could invoke all four conditions. Suppose he is a good cop that has never engaged in police brutality and he wishes that others wouldn't believe that he was a bad cop. Again, it is central to his conception of himself.

To answer these cases I turn back to the discussion of dispositional vulnerability from the previous section. The white participant in Robin DiAngelo's anti-racism training seminar, the white male police officer, they are not dispositionally vulnerable. Their sense of self-worth will be reaffirmed for them through structural features of our environment. Both need only turn on the TV or watch a movie to see countless positive portrayals of

[28] Thanks for Gabbrielle Johnson for raising this point.
[29] Thanks to Briana Toole for suggesting this case.

white males (and police officers). This is not something that exists for non-dominantly situated groups.

In general, this problem brings out another problem concerning the attempt to present necessary and sufficient conditions for when beliefs wrong. A challenge for accounts of norms on belief (or the epistemic processes that lead to the production of particular beliefs) is precisely this way in which standardly accepted norms of reasoning can be manipulated to wrong others. The police officer, the angry man in DiAngelo's seminar, they all attempt a move that Kristie Dotson (2018) refers to as a *seize of epistemic power*. Another term, to capture the way in which this move resembles emotional abuse and manipulation would be *epistemic manipulation*. For example, emotional abusers make use of norms that are part of a healthy relationship, and they use those norms to manipulate their partners. Part of why it is so difficult to recognize emotional manipulation as a form of abuse is how closely it resembles healthy behaviour. As the counterpart to emotional abuse and emotional manipulation, it can be similarly difficult to identify epistemic manipulation because of how closely it resembles epistemic norms we standardly accept as good epistemic norms. These very tools we can try to identify and articulate as ways of making sense of ways in which non-dominantly situated groups can be wronged by what others believe of them can be used and turned against them as tools of the oppressors. This is one of the many reasons I'm suspicious of attempts to give necessary and sufficient conditions for when beliefs wrong. I am generally suspicious of universal rules that apply in all contexts, environments and situations. Instead, as I'll suggest in the next chapter, perhaps there are general epistemic strategies we can employ when we recognize the moral stakes of a situation. However, recognizing the moral stakes, adjusting our epistemic behaviour to meet the stakes, etc., involves the development of a virtuous character that is sensitive to these issues as opposed to application of laws or rules. Employing the epistemic strategy alone will not be enough. The epistemic strategy must be applied with a certain amount of *moral sensitivity*. So again, we see, that the epistemic and moral cannot be kept entirely apart. It is a mistake to think that our moral lives and our epistemic lives are separate realms governed by conflicting obligations. To be moral we must not only do better, we must believe better as well.

To finally bring this chapter to a close, part of the point I have been trying to make is that there are many examples where we've been wronged by what someone believes of us that suggest some moral demand hasn't been met, and these examples all suggest that there is a pattern for determining when this is the case. However, the diversity of cases also suggests that there is no simple principle we can articulate. Although this is not an entirely satisfactory answer, as a general rule of thumb answers to difficult moral

problems are not going to be tidy like the solution to a math problem. If they are, then they're either solutions for some other kind of creature than a human being or solutions for a world that is quite different from the unjust one we live in. To fully understand our epistemic obligations with regard to what we believe of each other, we must understand what friendship requires of us, what being in committed relationships require of us, what being conscientiousness requires of us, etc. The final takeaway, however, that I hope to have made a case for is that fulfilling what we owe to others involves not only being sensitive to how we act towards each other and what we say of each other, but also what we believe of each other. There are moral demands on the beliefs we form and hold about each other.

CHAPTER III

The Moral Stakes of Our Beliefs

My goal in this dissertation has been to argue that even if racist beliefs were to reflect reality and thereby seem rationally justified, they are nonetheless wrong. Given the nebulous nature of determining which beliefs are racist, I have focussed on what I take to be the hardest case in which to see a moral wrong: beliefs that seem justified given the evidence. With regard to these beliefs, I have previously argued that the way in which they wrong is by neglecting an important moral difference between our beliefs about objects and natural events in the world and beliefs about other people. That is, how we morally ought relate to others affects how we ought respond to evidence concerning them before we settle belief. My first goal in this chapter is to expand upon the question of how we should respond to the evidence before settling belief. My other goal is to show that this moral requirement that emerges out of our relationships to others and in response to structural injustice reveals a moral dimension to the way in which our beliefs are *epistemically* justified. Accordingly, the question of what individuals should believe is, in part, a moral one.

To begin, let me introduce yet another case in which, as a result of the racist world we live in, the evidence seems stacked in favor of seemingly racist beliefs. Consider the following case from Gendler (2011) in which a social club's discriminatory membership practices have resulted in a stark racial divide between club members and staff members: only a fraction of the club's members are black, whereas all of the club's staff members are black.

> *Social Club.* Agnes and Esther are members of a swanky D.C. social club with a strict dress code of tuxedos for both male guests and staff members, and dresses for female guests and staff members. While preparing for their evening walk, the two women head toward the coat check to collect their coats. As they approach the coat check, they both look around for a staff member. As

> Agnes looks around she notices a well-dressed black man standing off to the side and tells Esther, "There's a staff member. We can give our coat check ticket to him."

In this case, it seems that Agnes's belief is justified. That the man standing in the lobby is black provides substantial evidential support for the belief that he is a staff member. Such a case, and others like it—recall that in *supposedly rational racist* we were assuming that that the diner is black provides substantial evidential support for Spencer's belief that the diner won't tip will—raise the troubling possibility that by simply attending to our epistemic obligations, we may contribute to the harms and wrongs of racism. On one hand, as we've seen discussed in Chapter 2, assuming that someone is a staff member on the basis of their race is a paradigmatic example of racism and is morally impermissible. On the other hand, the belief seems epistemically justified. If one had more hands, one might go further and add that it would be epistemically irrational *not to believe* that they are a staff member. After all, to refrain from believing in this case seems to involve a wilful disregard of the evidence.

In light of such cases, one might be tempted to follow Gendler (2011) and conclude that we face an irresolvable conflict between our moral obligations and what we epistemically ought to believe. I disagree. I suggest that these cases push us to rethink the assumptions that pit moral considerations and epistemic considerations against each other: namely, the evidentialist assumption that the only considerations relevant to whether you should believe are evidential considerations. That is, according to the evidentialist, moral considerations are neither here nor there when it comes to the question of whether a belief is justified given one's evidence.

In order to call into question whether Agnes or Esther are justified in believing what they believe on the basis of their evidence, we need first to understand what it is for your evidence to support believing *p*. In Section 3.1 I turn to looking closely at what evidentialism says, what it doesn't say, and to what degree various claims it makes are ones that are worth keeping in light of the troubling possibility that by simply attending to our epistemic obligations we may contribute to the harms and wrongs of racism. From this analysis, we will see that evidentialism leaves open both questions of how much evidence is enough evidence and the question of what determines how much evidence is enough evidence. As I'll show in Section 3.2 these open questions are where stakes-based encroachment accounts in epistemology gain traction: non-evidential considerations plausibly play a role in settling the question of how much evidence is enough evidence.

I am not the first to note this point. Defenders of pragmatic encroachment in epistemology have argued that non-paradigmatically epistemic features of a subject's environ-

ment can make a difference to what that subject is justified in believing or what they know. Consider the following example, if you are deathly allergic to nuts, the threshold your evidence must pass for you to be justified in believing that a dish does not contain nuts is higher than for someone lacking such an allergy. In Sections 3.2.1-3.2.6 I canvas various accounts of pragmatic encroachment that have been offered in the literature and show how they challenge and weaken the evidentialist thesis. In so doing, I will suggest that the motivating cases for pragmatic encroachment will also help us motivate *moral encroachment*. Whereas traditional stakes-based encroachment accounts are articulated in terms of *practical* stakes, such accounts leave open the question of *which* non-paradigmatically epistemic features of a subject's environment can make a difference to what that subject is justified in believing or what they know. This leaves open the possibility that features that matter *morally* can affect whether a subject is justified or whether a subject knows. Similar arguments can be found in Fritz (2018), Pace (2011), Moss (2018), and Schroeder (2018b) and I will expand upon these ways in which others have similarly motivated moral encroachment in Section 3.3.

Once the stage is set for an account of *moral encroachment*, in Section 3.4 I address three questions that arise for such an account. First, what would the low-stakes variants of the traditional pragmatic cases be; second, what is it that makes these cases high stakes; and third, is what makes these cases high stakes the same in both the pragmatic and the moral domains. I argue that our epistemic practices are constrained by a kind of moral sensitivity. This moral sensitivity amounts to not only a mere recognition of an unjust environment, but also an adjustment of one's threshold for justification to meet the moral stakes of one's environment. Finally, in Section 3.5 I consider three objections to moral encroachment. First, that moral encroachment is an incomplete account of the ways in which beliefs can wrong, i.e., of doxastic morality. Second, that moral encroachment has perverse consequences. Third, that moral encroachment is too demanding.

3.1 Rethinking the Evidentialist Assumptions

Let us start with another case to get the intuitive feel for evidentialism. Suppose your friend Veronica tells you that she knows the identity of the Zodiac Killer. Intrigued, you follow up. It turns out that Veronica believes that Ted Cruz is the Zodiac Killer. Not wanting to dismiss her out of hand, presumably you'd go on to ask her for what evidence she has in support of her belief. If you found out that she *hopes* that Ted Cruz is the Zodiac Killer, because it would be funny, that would be irrelevant to the question of whether in fact Ted Cruz is the Zodiac Killer. After all, that Veronica hopes Ted Cruz is the Zodiac

Killer has no bearing on whether or not Ted Cruz is the Zodiac Killer.

The purpose of this vignette is to illustrate the simple point that spelling out the relationship between evidence and justification is not a straightforward matter. This point complicates the standard evidentialist picture since, although the evidentialist is correct to note that there is a connection between justification and evidence and that we should believe in accordance with the evidence, this picture leaves open two distinct questions about what settling *in accordance with* the question amounts to. How much evidence is enough evidence? What determines how much evidence is enough evidence? In what follows, I use this example of Veronica to demonstrate that spelling out what it is to believe in accordance with the evidence first requires settling these two questions. Further, these questions are sensitive to non-evidential considerations.

Returning to Veronica, she might have *some* evidence that Ted Cruz is the Zodiac Killer. For example, Ted Cruz somewhat resembles the police sketch of the Zodiac Killer. This evidence, however, is not *sufficient* to make Veronica's belief that Ted Cruz is the Zodiac Killer justified.[1] For the belief to be justified, she would need more evidence. This simple case illustrates that when we try to determine whether a belief is justified, we look not only for evidence—reliable signs, symptoms, or marks of that which it is evidence of, e.g., reasons that raise the probability of p being true—we also consider the *strength* of the evidence. In short, whether a belief is justified is a matter of both what the evidence is and second, whether the evidence is sufficient to justify the belief. We can articulate this minimal intuitive gloss on evidentialism as follows.

> *Support*. For each proposition p and body of evidence E, there is some threshold for justification t, such that it is permissible to believe p on the basis of E only if the $\text{pr}(p|E) \geq t$.

Here pr(p|E) corresponds to some way of understanding the *strength* of the evidence E for the proposition under consideration, and the threshold t corresponds to a *sufficiency* condition the strength of one's evidence must either meet or surpass to be justified. For example, consider an easy case, such as whether you are justified in believing that the kettle has boiled. The sound of a kettle whistling raises the probability of it being true that the kettle has boiled. Further, it raises the probability higher than if you had just glanced at your watch and seen that the kettle has been on for two minutes. Both make it more likely than not that the kettle has boiled, but only the sound of the whistling raises the

[1] Stepping away from cases such as the identity of the Zodiac Killer, we see this constraint on our beliefs everywhere in our daily life. If you are trying to decide whether to pack an umbrella, you might look outside or check the weather report to gather evidence before settling your belief. Your friend, on the other hand, who believes that it will rain because they *want* it to rain believes irrationally.

probability to a degree that is sufficient for being justified in believing that the kettle has boiled. However, that threshold need not be the same for every belief. Presumably, in some circumstances you would want stronger evidence, i.e., you would want the threshold to be higher. In others, you might only require weaker evidence, i.e., for the threshold to be lower. For example, in adjudicating someone's guilt in a criminal case, the threshold is higher than when adjudicating a playground dispute. All of this is consistent with thinking that there is some threshold that the strength of one's evidence must cross in order to justify belief.

In line with this minimal gloss on the evidentialist intuition, in Section 3.4 I argue that many factors, including the moral demands of our environment, affect where we place the threshold. For example, if you are in a morally demanding situation, e.g., a context in which you risk contributing to the harms of racism given what you believe, the threshold is higher than when you are not in such situations, e.g., when you are deciding whether to bring an umbrella to work. However, many evidentialists reject the idea that the threshold for justification can be *shifty* in this way because it conflicts with some more restrictive accounts of evidentialism.

Many evidentialists are not satisfied with only *Support*. This dissatisfaction leads to the endorsement of restrictions on justification that are designed to restrict the shiftiness of the threshold. One commonly endorsed principle that is considered central to evidentialism is the following.

> *Supervenience.* For any two subjects S and S', necessarily, if S and S' have the same evidence for/against *p*, then S is justified in believing that *p* iff S' is, too.[2]

Supervenience is designed to capture the thought that, for example, two jurors charged with adjudicating someone's guilt shouldn't have differing thresholds for justification. If they both have the same evidence for a guilty verdict, then if one is justified in believing that the defendant is guilty, so too is the other juror.[3]

A further restrictive version of evidentialism suggests that there is only ever one belief that is justified given the evidence—*Uniqueness*.

[2] See Fantl and McGrath (2002) and Conee and Feldman (2004). Fantl and McGrath call this principle 'purism'.

[3] Note that the subjective Bayesian would reject this formulation of *Supervenience*, but they would accept the following version: for any two subjects S and S', if S's evidence and S"s evidence supports *p* to the same degree, if S is justified, S' is too. Note that the difference is that is not merely a matter of having the same evidence, the evidence must support the belief to the same degree for both agents. Pragmatic encroachment is usually taken to be incompatible with both. The account of moral encroachment that I will develop in this chapter gives us counterexamples to this second form of *Supervenience*, but not to the first. In part because of the difficulties of creating cases contrasting pairs of high-low cases.

> *Uniqueness.* 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. [4]

Returning to the cases that motivated us at the beginning of the chapter, according to the restrictive version of evidentialism that is jointly committed to *Support* and *Supervenience*, anyone who finds themselves in the same situation as Agnes and Esther is similarly justified in believing that the black man standing in the lobby is a staff member. That is, if Agnes and Esther are justified, so too are they; and if Agnes and Esther are not justified, neither are they. If we add *Uniqueness*, then that is the only belief that is justified given the evidence. That is, it is not only the case that one ought believe that he is a staff member, but also the rest of us are irrational if we don't believe that he is a staff member.

Although both *Supervenience* and *Uniqueness* have been taken to be central components of evidentialism, both are stronger than the initial gloss we gave of evidentialism in terms of *Support*. It is open for debate whether our intuitive practices of requiring evidence in support of our beliefs requires all three theses—*Support*, *Supervenience*, and *Uniqueness*. I suggest that the seemingly irresolvable nature of the conflict at the heart of this dissertation—the conflict between what we ought morally to believe and what we ought epistemically to believe—only arises because of the more restrictive accounts of evidentialism masquerading as intuitive glosses of the more general evidentialist intuition that one ought believe in accordance with the evidence.

An important upshot of the account I develop in this chapter is that although stakes-based encroachment views are often considered rejections of evidentialism, they can remain compatible with the most intuitive form of evidentialism. To embrace moral encroachment, then, does not require denying an important connection between justification and evidence. All that moral encroachment asks for is the recognition that what counts as *enough* evidence to justify belief can vary according to practical factors *and* the moral stakes of the agent's situation.

Moving forward my strategy is to show how we can resolve this conflict between our moral and our epistemic obligations while preserving the basic evidentialist intuition understood as a commitment to *Support*.

As a matter of framing, you might respond that this is not so much an issue for eviden-

[4] From Feldman (2007, pp. 205). Alternatively, White (2005, pp. 445) presents uniqueness as follows: "Given one's total evidence, there is a unique rational doxastic attitude that one can take to any proposition." Further, as before with *Supervenience*, there is also a weaker version of *Uniqueness* which restricts impermissivism to just one person at a particular time. That is, for each agent at each time, the evidence justifies only one proposition.

tialism, rather it is a challenge for the moral requirement to occupy the moral standpoint with regard to what we believe of each other. That this requirement conflicts with a seemingly orthodox view of belief and justification of belief is a strong reason against thinking that there could be such a requirement. Thus, the burden is not on evidentialism to accommodate this moral-epistemic requirement.

Rather, the burden is on those that defend the distinction to do one of five things. First, we could explain away the intuitive plausibility of evidentialism insofar as it seems responsible for generating this kind of conflict. Alternatively, we could show that there is in fact no conflict between the moral and epistemic ought by rethinking and questioning the assumptions that seem to pit them together (as I will be suggesting). Third, perhaps we could show that although there is a conflict between epistemic and moral demands, there is a way of adjudicating between them such that we can say there is something Spencer and Agnes ought all-things-considered believe. Fourth, we could deny that there's anything that Spencer or Agnes ought all-things-considered believe while *also denying* that that means we must resort to the pessimistic conclusion that this conflict is irresolvable. That is, when oughts seem to conflict there is no conflict, there is one sense in which you should do x and *a different sense* in which you should do y. Finally, fifth, we could just endorse the pessimistic conclusion that there is no way to determine what to do in cases of conflicts.

My goal in this chapter is the second approach—show that there is in fact no conflict between the moral and epistemic ought by rethinking and questioning the assumptions that seem to pit them together. I leave the task of canvassing reasons for preferring this approach to the third, fourth, and fifth approach for Chapter 4. With regard to the first alternative, I will briefly note that although it might be tempting given the many challenges to evidentialism despite it being considered the orthodoxy, the standard ways of challenging the intuitive plausibility of evidentialism will not help us in this case. For certain classes of belief, evidentialism fails as an explanation of their justificatory status. For example, reasons to believe mathematical truths or a priori truths do not rest on one's evidence. Nor do cases of self-knowledge require that an agent believe in accordance with her evidence. For empirical beliefs, on the other hand, like your friend's belief that Ted Cruz is the Zodiac Killer, evidence does seem like what we look for to determine whether the belief is rational and justified. Beliefs about others are also empirical in this way. So, there is reason to think that evidentialism is a requirement on such beliefs. My goal, then, is to show that the requirement to occupy the moral standpoint, and for moral considerations to affect the *epistemic* justification of our beliefs, is not at odds with a weak version of evidentialism. To do so, we find friends amongst the pragmatic encroachers.

3.2 Varieties of Pragmatic Encroachment

Suppose you are not allergic to nuts but simply have a general dislike for them. In such a case, a waiter's assurance that they're *pretty sure* there aren't nuts in any of the dishes would be sufficient for you to justifiably believe that there aren't nuts. However, if your dining partner Sarai is deathly allergic to nuts, then the waiter's assurance no longer suffices. Sarai, given the stakes, should not believe that none of the dishes contain nuts. Both you and Sarai, however, have the very same evidence that supports the belief *none of the dishes contain nuts* to the very same degree. That is, if we suppose the evidence—the waiter's testimony—supports that the dish doesn't contain nuts to degree .7 for you, it also supports that the dish doesn't contain nuts to degree .7 for Sarai. However, whereas you are justified in believing that none of the dishes contain nuts, we tend to be resistant to the claim that Sarai is similarly justified in believing that none of the dishes contain nuts. As I have been pressing, the evidential situation for both you and Sarai is identical. But, there is a difference in our intuitive judgements concerning justification. This suggests that whatever difference there is between our judgments about justification must come down to something other than the evidence. By simply changing a practical feature of the scenario—what is at stake for Sarai—we have also thereby changed what she is epistemically justified in believing. Although the evidentialist might reject practical considerations as irrelevant to the question of justification, our intuitions suggest otherwise.

Defenders of pragmatic encroachment accounts in epistemology use cases like the one above to argue that practical features of one's environment, not evidence alone, can change the epistemic status of a believer—e.g., one can change from being justified to not being justified, from having to knowledge to lacking knowledge. Pragmatic encroachment, despite having historical roots in Locke (1975)/1689, is often considered a relatively recent philosophical thesis that offers an invariantist alternative to contextualist treatments of 'knows'.[5] Understood in this way, it can initially be difficult to see how pragmatic encroachment might offer a structural analogue to the moral challenge I raise against evidentialism. The challenge I raised concerns the rational permissibility of belief formation in light of moral concerns. But if pragmatic encroachment only concerns the semantics

[5]From Grimm (2016):

> According to Locke, for instance, "it is very wrong and irrational way of proceeding, to venture a greater good for a less, upon uncertain guesses and before a due examination be made, proportionable to the weightiness of the matter, and the concernment it is to us not to mistake" (2.21.66); and again, "where the assent one way or the other is of no importance to the interest of anyone ... there 'tis not strange that the mind should give itself up to the common opinion, or render itself to the first comer' (1975: 717)

of 'knows' and the question of whether knowledge ascriptions are context-sensitive, what does that have to do with the rationality of belief? As I'll show, by canvassing the various versions of pragmatic encroachment, pragmatic encroachment is not merely about knowledge and knowledge ascriptions.

To start, our concept of knowledge involves various components that can be given a probabilistic gloss. For example, knowledge entails justified belief. When you know that *p*, that is partly constituted by your believing that *p*. So, when we attribute knowledge to someone we are also attributing a belief to them. And when we attribute a belief to someone, we can ask how sure, how confident, and how strongly they believe that *p*. Similarly, for that belief to count as knowledge it must be justified. But as we saw in the discussion of evidentialism, we can ask how much evidential support is required to justify belief. In light of this, as I have been suggesting, it is a misconception to think that pragmatic encroachment only concerns knowledge.[6] As I'll show, some pragmatic encroachment proposals argue that pragmatic considerations do not apply directly to knowledge, but rather they enter into our assessments of knowledge through pragmatic constraints on belief or pragmatic constraints on justification. Whereas other defenses of pragmatic encroachment emerge from a muddying of the distinction between practical and theoretical rationality with regard to how to balance the competing goals of believing the truth and avoiding error, and further questioning whether epistemic rationality is exhausted by evidential considerations. The phenomenon of pragmatic encroachment raises the general question of whether there are any epistemic notions that aren't sensitive to an agent's interests.

Following the presentation of the varieties of pragmatic encroachment, I then turn to sketching the most innocuous and least-committal version of pragmatic encroachment I can. Doing so allows us to explore the challenge pragmatic encroachment raises for evidentialism while taking on as few other commitments as possible and thus being neutral between the various theories. Although still controversial, my hope is that by presenting the least-committal version of the basic pragmatic encroachment intuition we can see the shape of the challenge pragmatic encroachment offers against evidentialism, and how it modifies evidentialism in light of the response. And then in Section 3.3 we can explore if the same strategy can be extended to accommodate moral considerations. This way, we need not commit ourselves to any particular theory of pragmatic encroachment in presenting a story about moral encroachment.

[6] We can also see this limited version of pragmatic encroachment expressed in the *Stanford Encyclopedia of Philosophy* definition of pragmatic encroachment as merely the thesis that a difference in pragmatic circumstances can constitute a difference in knowledge.

3.2.1 The Contextualist Roots of Pragmatic Encroachment

As noted earlier, pragmatic encroachment, despite having historical roots in Locke, is often considered a relatively recently philosophical thesis that offers an invariantist alternative to contextualist treatments of 'knows'. So to understand the pragmatic encroachment thesis, it will be useful to start with the view that it is often contrasted against: contextualism. Contextualism about 'knows' is motivated by the following pair of cases from DeRose (1992).

> *Bank Case A.* My wife and I are driving home on a Friday afternoon. We plan to stop at the bank on the way home to deposit our paychecks. But as we drive past the bank, we notice that the lines inside are very long, as they often are on Friday afternoons. Although we generally like to deposit our paychecks as soon as possible, it is not especially important in this case that they be deposited right away, so I suggest that we drive straight home and deposit our paychecks on Saturday morning. My wife says, "Maybe the bank won't be open tomorrow. Lots of banks are closed on Saturdays." I reply, "No, I know it'll be open, I was just there two ago on Saturday. It's open until noon."

> *Bank Case B.* My wife and I drive past the bank on a Friday afternoon, as in Case A, and notice the long lines. I again suggest that we deposit our paychecks on Saturday morning, explaining that I was at the bank on Saturday morning only two weeks ago and discovered that it was open until noon. But in this case, we have just written a very large and important check. If our paychecks are not deposited into our checking account before Monday morning, the important check we wrote will bounce, leaving us in a *very* bad situation. And, of course, the bank is not open on Sunday. My wife reminds me of these facts. She then says, "Banks do change their hours. Do you know the bank will be open tomorrow?" Remaining as confident as I was before that the bank will be open then, still, I reply, "Well, no. I'd better go in and make sure."

That I know that the bank will be open on Saturday in *Bank Case A* seems true, but I also seem to be saying something true in *Bank Case B* when I concede that I don't know that the bank will be open on Saturday. However, the evidence supports my belief to the same degree, and it is not as though I am in any better epistemic position in A than in B, so following *Supervenience*, if I know in *Bank Case A*, I must also know in *Bank Case B*.

In light of these cases, epistemic contextualists argue that the content of 'S knows that p' depends upon the context of utterance (i.e. the attributer's context). A purported benefit to this explanation is that it does not undermine the following thesis.

> *Intellectualism.* The factors in virtue of which a true belief amounts to knowledge are exclusively truth-conducive, in the sense that their existence makes the belief more likely to be true.

In the discussion of evidentialism from Section 3.1 we captured something similar to *Intellectualism* in our definition of evidence as that which raises the probability p of being true. We can see the evidentialist as committed to a more restricted version of intellectualism in which the factors in virtue of which a true belief amounts to knowledge is determined by the strength of the evidence alone. After all, the evidentialist defines evidence as anything that raises the probability of p being true. So, the evidentialist take on *Intellectualism* would be that the factors in virtue of which a true belief amounts to knowledge are exclusively evidential.

Hawthorne (2003), however, argues that an unattractive feature of contextualism is that it cannot maintain the intuitive connections between knowledge, assertion, and practical reasoning. In response, he develops an account of pragmatic encroachment as an invariantist alternative to contextualism. Relatedly, Stanley (2005) also argues that contextualism is counter-intuitive and that the counter-intuitive results that it delivers on variants of the bank cases also displays the way in which *Intellectualism* stands in conflict with an intuitive connection between knowledge and action. I turn now to presenting in more detail Stanley (2005)'s account of pragmatic encroachment. At the same time that Stanley was developing his view, Fantl and McGrath (2002, 2007, 2009) over a series of papers and a book were developing a similar account but from different motivations. What remains central to all three accounts, and later the combined Hawthorne and Stanley (2008) account, is the importance of preserving the intuitive connection between knowledge and action.

3.2.2 Explaining Encroachment: Knowledge and Action

Stanley (2005)'s motivating thought is that our intuitions about whether someone knows that p seem to vary according to non-epistemic features of that person's situation. For example, if someone is in a risky situation we feel that she must be in a particularly good epistemic position. Whereas we are more permissive when she is not in a risky situation. Recall that when not much was at stake, it is true that I know that the bank will be open on Saturday. However, when I'm in a riskier situation with regard to my

finances, it no longer seems that I know that the bank will be open on Saturday. Further, not only do our intuitions change based on the risk, they also seem to change based on what it is salient to either the knower or ourselves. Stanley argues that in a variation of *Bank Case B* in which I am ignorant of the high stakes if the bank's closed and I don't get a chance to deposit my paycheck, 'I know the bank will be open on Saturday' is not true, whereas according to the contextualist it would be true.

In short, Stanley's argument for pragmatic encroachment proceeds by creating similar low-stakes/high-stakes pairs of cases to the Bank Cases that motivated contextualism and showing how our intuitive responses suggest an alternative to contextualism. The basic formula for creating these cases is to first take an agent and put them in a scenario where it would be reasonable to think that they know. The contrasting case, then involves a hypothetical parallel in which the cost of being wrong no longer makes it reasonable to think that they know. We can also see how the formula for creating cases is a direct challenge to the *Supervenience* component of evidentialism. Given *Supervenience*, our judgements about knowledge should not shift merely because of a difference with regard to the cost of being wrong. It is worth making this formula explicit because it is common to all accounts of pragmatic encroachment. Distinctive of Stanley's approach, however, is the weight he places on our intuitive responses to these cases.

Stanley suggests that what our intuitions are tracking in such cases is a conceptual link between knowledge and action. That is, to know that p is related to being rationally permitted to act upon the belief that p. He suggests that given that the latter is a function of practical facts, it wouldn't be a surprise if the former was as well. His suggestion, then, is that given this connection with action, knowledge is an interest-relative notion. That is, whether someone knows that p at a time t depends at least in part upon practical facts. Although his argument proceeds through a polling of intuitions, he defends this approach by noting the following.

> The role of these intuitions is not akin to the role of observational data for scientific theory. The intuitions are instead intended to reveal the powerful intuitive sway of the thesis that knowledge is the basis of action. Someone who denies that we have many of these intuitions is denying the pull of the link between knowledge to action. (pp. 12)

So, what explains the intuitions we have in the various variants of the bank case that Stanley presents is a plausible conceptual connection between knowledge and action? To capture these intuitions, Stanley puts forward the following view as an alternative to contextualism.

> *Interest-relative Invariantism.* Whether or not someone knows that *p* may be determined in part by practical facts about the subject's environment. (pp. 86)

In so doing, Stanley straightforwardly rejects intellectualism and instead endorses the view that whether a belief constitutes knowledge depends in part on non-truth-conducive practical matters such as the cost of being wrong.

Now, despite the discussion thus far being couched in terms of knowledge it can be extended to evidence given that both Stanley and Hawthorne endorse Williamson (2000)'s claim that E=K, that is S's evidence is S's knowledge. What one knows can be identified with a person's total evidence. So, in the last pages of his book, Stanley begins to question whether all epistemic notions are interest-relative. After all, if knowledge is interest-relative and knowledge just is an agent's total evidence, perhaps evidence is interest-relative as well. He asks,

> If all epistemic notions are interest-relative, then evidence is interest-relative as well. If evidence shares the interest-relativity of knowledge, then two people who do not share the same practical situation will not in general have the same evidence. (pp. 124)

We see that pragmatic encroachment is not only a theory about knowledge, it is also a theory about evidence, and in general about the boundaries of what counts as 'epistemic'. To gesture ahead, as we'll see in discussion of Schroeder (2012a)'s account of pragmatic encroachment, if we think that the only reasons that determine the rational permissibility of belief are epistemic reasons, given the phenomena that motivates pragmatic encroachment we cannot presume that that claim is identical with the claim that the only evidence (or strength of evidence alone) determines the rational permissibility of belief. So, we see another way in which pragmatic encroachment challenges traditional evidentialist theses.

But, before we turn to Schroeder (2012a)'s account of pragmatic encroachment, we must first consider an account of pragmatic encroachment more similar in spirit to that of Stanley. That is, an account similarly motivated by a desire to preserve the relation between knowledge and action, a distrust of *Intellectualism*, and a rejection of *Supervenience*. This is the account offered by Fantl and McGrath (2002, 2009).

3.2.3 Explaining Encroachment: Knowledge-Justification and Action

Fantl and McGrath (2002) start with a familiar pair of cases, this time concerning trains.

Train Case 1. You're at Back Bay Station in Boston preparing to take the commuter rail to Providence. You're going to see friends. It will be a relaxing vacation. You've been in a rather boring conversation with a guy standing beside you. He, too, is going to visit friends in Providence. As the train rolls into the station, you continue the conversation by asking, "Does this train make all those little stops, in Foxboro, Attleboro, etc." It doesn't matter much to you whether the train is the "Express" or not, though you'd mildly prefer it was. He answers, "Yeah, this one makes all those little stops. They told me when I bought the ticket." Nothing about him seems particularly untrustworthy. You believe what he says.

Train Case 2. You absolutely need to be in Foxboro, the sooner the better. Your career depends on it. You've got tickets for a southbound train that leaves in two hours and gets into Foxboro in the nick of time. You overhear a conversation like that in Train Case 1 concerning the train that just rolled into the station and leaves in 15 minutes. You think, "That guy's information might be wrong. What's it to him whether the train stops in Foxboro? Maybe the ticket-seller misunderstood his question. Maybe he misunderstood the answer. Who knows when he bought the ticket? I don't want to be wrong about this. I'd better go check it out myself."

As before in the bank cases, in the first case we have a different judgement than in the second case. When not much is at stake, you have good enough evidence to know that the train stops in Foxboro. However, when more is at stake, you do not have good enough evidence to know that the train stops in Foxboro. Fantl and McGrath (2002) use this case to suggest that epistemic justification—the kind of justification that entails knowledge—is not simply a matter of the evidence that one has. In both train cases, you're supposed to have the same evidence. But in one, you are justified. And in the other, you are not. So, not only do we have a violation of *Supervenience* (which Fantl and McGrath (2009) refer to as *Purism*), we also see the way in which the discussion concerning pragmatic encroachment is not only about knowledge, it is also about justification.

In their paper, Fantl and McGrath (2002) use our judgment in this case to defend the following pragmatic condition on epistemic justification.

(PC) S is justified in believing that p only if it is rational for S to prefer as if P.

This principle follows from the following principle connecting knowledge to action. Suppose S knows that p—e.g., the door is locked—and assume that S knows that if p, then A is

the thing to do—e.g., if the door is locked, don't double back to check the door is locked. If I know that p, and if I know that if p, then A, then I'm rational in doing A. The train cases illustrate that knowledge shifts because what it is rational to do shifts.

Later, in their book, Fantl and McGrath (2009) build their account of pragmatic encroachment is around the following principle, Knowledge-Justification.

> (KJ) If you know that p, then p is warranted enough to justify you in ϕ-ing, for any ϕ.

But, why suppose that (KJ) is true? On similar grounds to their earlier paper and to Stanley and Hawthorne, Fantl and McGrath argue that (KJ) makes sense of our habits of citing knowledge to criticize or justify action. When you know that p, you can count it in action as well as in belief.[7] Whereas Fantl and McGrath (2002) proceed with some intuition marshalling, Fantl and McGrath (2009)'s argument for pragmatic encroachment is supposed to follow directly from our understanding of knowledge (and the related concept of justification) to action.

They start with the following two intuitive theses about knowledge.

> (Fallibilism) You can know that p even if there is a non-zero epistemic chance for you that not-p.

> (Purism) If two subjects are just alike in their strength of epistemic position with respect to p, then they are just alike with respect to whether they are in a position to know that p.[8]

[7] Similarly, Hawthorne and Stanley (2008) agree. They note that our our ordinary folk appraisals of the behaviour of others suggest that the concept of knowledge is intimately intertwined with the rationality of action, and they suggest the following principle connecting action to knowledge.

> *The Action-Knowledge Principle.* Treat the proposition that p as a reason for acting only if you know that p.

This principle is offered as a way of making sense of our use of 'know' to criticize the actions of others. Further, in their defence of this principle, they also note that

> an analogous principle seems plausible for reasons for belief, viz.: Treat the proposition that p as a reason for believing q only if one knows that p. One attractive feature of the Action-Knowledge Principle is that it unifies the practical and theoretical domain of reasons: if it is correct, then proper reasons for belief and reasons for action have a uniform nature.

So we see the endorsement of something similar to Fantl and McGrath's unity thesis which says that if p is warranted enough to be a reason you have to believe, p is also warranted enough to be a reason you have to act (or anything else).

[8] We have been calling this thesis *Supervenience*. Also, implicit in Purism/Supervenience is a commitment to intellectualism. As Fantl and McGrath explain, when the Purist talks of "strength of epistemic position" they only mean truth-relevant considerations. How much is at stake for you in whether p is true,

Purism does have an intuitive appeal because whether you know is a matter of your relation to the truth. Should purism turn out to be false, then an unattractive consequence would be that you could lose knowledge not by losing information or forgetting something, but simply because of a change in your stakes. However, we cannot preserve both purism and fallibilism. The price of purism is giving up fallibilism. But fallibilism is also compelling because it seems to be required in order for us to know much of anything important and because many of things it seems, intuitively, we know are such that there is a non-zero chance that they are false. Fantl and McGrath suggest that of these two theses, it is purism that must go, and it follows from the intuitive plausibility of KJ and Fallibilism. Their argument goes as follows.

Let's start by assuming that fallibilism is true. Now suppose that I fallibly know p where p is that the door is locked. Following KJ, p is warranted enough to justify me in not doubling back to check that the door is locked. But, because of fallibilism, there's a chance that my door is not locked. So, there's a hypothetical case in which *that the door is locked* isn't warranted enough to justify me in not doubling back to make sure the door is locked. In the hypothetical case, I must double back. So, in the hypothetical case I'm not in a position to know that the door is locked, even though in the hypothetical case I'm in the same strength of epistemic position—alternatively, I have the same strength of evidence—with respect to the proposition that the door is locked as I am in the actual case in which I know the door is locked. Therefore, purism must be false.

This is very quick sketch, but we can see the general picture of the approach and can again set aside the misconception that pragmatic encroachment is only about knowledge. Knowledge is practically sensitive, according to Fantl and McGrath, because justification is practically sensitive. So again, we see a broadening of interest-relativity to the epistemic concepts related to knowledge, and not only to knowledge itself. This again, should be unsurprising given that knowledge entails justified belief.[9]

3.2.4 Explaining Encroachment: Belief and Action

In contrast with the accounts presented above, Ganson (2008), Weatherson (2005), and Ross and Schroeder (2012) argue that pragmatic considerations get their foot in the door in virtue of a pragmatic constraint on belief.

---

whether p would make you happy, etc. are not part of the strength of your epistemic position. If we again understand evidence for p as broadly meaning anything that raises the probability that p is true, then we can substitute 'strength of evidence' for 'strength of epistemic position'. But, as we'll see in Schroeder (2012a), there is reason to be suspicious of such a substitution.

[9]Further, Fantl and McGrath (2009) extend this general framework to argue that there's pragmatic encroachment on justification. That is, we can argue that purism is not only false for knowledge ascriptions, but so for justification.

The arguments for pragmatic encroachment presented by Hawthorne, Stanley, and Fantl and McGrath are presented in the language of belief and disbelief. As such, they do not take into account credences or degrees of belief. When we take into account the difference between outright belief and degrees of belief, Ganson (2008), Weatherson (2005), and Ross and Schroeder (2012), despite disagreeing about the precise nature of belief, argue that there is no pragmatic condition on justification, rather, there is a pragmatic condition on belief.

To expand, knowledge-action accounts, as their name suggest, posit connections between knowledge and action. The belief-action accounts, however, argue that pragmatic considerations are relevant only to the question of whether acting as if *p* is reasonable, so that means pragmatic considerations are relevant to whether the agent is willing to act as if *p*. The question of whether an agent is willing to act as if *p* is one that's settled by whether she has an outright belief. So, there's a pragmatic constraint on the justification of outright believing that *p*. Weatherson (2005) uses this argument to suggest that there is no pragmatic encroachment in *epistemology*, rather, pragmatic encroachment is a thesis for philosophy of mind. He notes,

> interests matter not because they affect the degree of confidence that an agent can reasonably have in a proposition's truth. (That is, not because they matter to epistemology,) Rather, interests matter because they affect whether those reasonable degrees of confidence amount to belief. (That is, because they matter to philosophy of mind.) There is no reason here to let pragmatic concerns into epistemology. (pp. 435-6)

Weatherson (2012), however, changes his mind and now argues that the interest-relativity of knowledge ascriptions goes beyond the interest-relativity of corresponding belief ascriptions. He motivates this change in light of the important role that knowledge plays in decision theory, and thus can also be classed in the same group as Hawthorne, Stanley, and Fantl and McGrath given the connection between decision theory and action. The difference between Weatherson and Fantl and McGrath is that whereas Fantl and McGrath argue that knowledge is practically sensitive because the justification condition, or the rational permissibility condition, on belief is practically sensitive, Weatherson argues that practical considerations impact knowledge directly.

The takeaway, however, is simply that we can see here again that pragmatic encroachment is not merely a thesis about knowledge. Rather, it can be construed more broadly as the thesis that non-evidential considerations, such as one's practical interests or one's willingness to act on p, determine the rationality of or justificatory status of one's belief.

3.2.5 Explaining Encroachment: Epistemic $\neq$ Evidential

As I have been stressing, construing pragmatic encroachment as merely a thesis about knowledge ascriptions misses the complexity of the challenges and the questions raised by the motivating intuition guiding the various pragmatic encroachment proposals offer. One such question is whether by epistemic considerations we mean only evidential considerations. As noted in my representation of evidentialism, it is plausible to think of evidence for *p* as anything that raises the probability that *p* is true. Further, if we consider epistemic rationality is the kind of rationality entailed by knowledge and thus constrained by truth-conducive reasons, then perhaps epistemic considerations are merely evidential considerations.

The motivating intuition guiding the various pragmatic encroachment proposals, however, calls this assumption into question. If practical interests can affect knowledge, and if epistemic rationality is the kind of rationality that is entailed by knowledge, then practical interests can affect epistemic rationality. But epistemic rationality is more than just about knowledge, it is about what it is rational and reasonable to believe. So, we see encroachment encroaching deeper. Practical interests can affect what it is reasonable to believe. So, when it comes to the rational permissibility of belief, that is determined by more than strength of evidence alone. Thus, we again seem to challenge the evidentialist assumption that only strength of evidence determines the rational permissibility of belief.

This line of argument is presented by Schroeder (2012a). Schroeder, following Stanley and Fantl and McGrath, starts with a pair of cases modelled off of the earlier bank case from DeRose. However, he removes the knowledge ascriptions to focus on our judgements about what it is reasonable to believe.

> *Bank-Low.* Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their pay checks. It is not important that they do so, as they have no impending bills. But as they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoons. Hannah remembers the bank being open on Saturday morning a few weeks ago, so she says, 'Fortunately, it will be open tomorrow, so we can just come back.' In fact, Hannah is right—the bank will be open on Saturday.

> *Bank-High.* Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their pay checks. Since their mortgage payment is due on Sunday, they have very little

> in their account, and they are on the brink of foreclosure, it is very important that they deposit their pay checks by Saturday. But as they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoons. Hannah remembers the bank being open on Saturday morning a few weeks ago, so she says, 'Fortunately, it will be open tomorrow, so we can just come back.' In fact, Hannah is right—the bank will be open on Saturday.

Whereas it seems reasonable to believe in the first case, it does not seem reasonable to believe in the second. To see how this could be the case, consider the following analogue to *Support* that Schroeder (2012a, pp. 268) offers.

> *Support**. It is epistemically rational for S to believe *p*, just in case *p* is adequately supported by S's evidence.

Both *Support* and *Support** leave open the question of what determines the threshold for adequate support. That is, when is the evidence such that a belief is adequately supported, and what determines when the evidence is strong enough to adequately support *p*? It is here, in this question of what it takes for a proposition to be adequately supported by the evidence, that non-evidential considerations come into play.

Schroeder argues that in addition to reasons to believe and reasons to not believe, our intuitions in the cases that motivate pragmatic encroachment suggest another class of reasons: reasons to withhold. When we look at the cases used to motivate pragmatic encroachment, we see that they tend to involve some kind of disadvantage that might arise in forming the belief, in making up your mind. They involve reasons to withhold. Costs of error are reasons to withhold and these are epistemic reasons that are not exhausted by the evidence. Why?

> Reasons to withhold cannot be evidence because evidence is exhausted by evidence which supports p and evidence which supports ¬p. But the evidence which supports p is a reason to believe p, and the evidence which supports ¬p is reason to believe ¬p. Consequently the reasons to withhold must come from somewhere else. So they cannot be evidence. (pp. 276)

Finally, it is also here that we see a suggestion for how the pragmatic encroachment challenge can be seen as not being in conflict with a version of weak evidentialism. As Schroeder (2012a, pp. 282) suggests,

> all of this is consistent with the thesis that it is epistemically rational for S to believe p just in case S has adequate evidence for p. Reasons to withhold

> simply raise the bar on how good the evidence needs to be, in order to be adequate.

On one final note, pragmatic encroachment not only gains its plausibility from these intuitive cases that conflict with the prescriptions of evidentialism, it also gains plausibility with regard to how it accommodates the two goals we attempt to balance when forming beliefs—believing truly and avoiding error. As James (1896) notes, our beliefs are governed by two intellectual goals that often pull in two different directions.

> Believe truth! Shun error!—these, we see, are two materially different laws; and by choosing between them we may end by coloring differently our whole intellectual life. We may regard the chase for truth as paramount, and the avoidance of error as secondary; or we may, on the other hand, treat the avoidance of error as more imperative, and let truth take its chance.[10]

On the one hand we try to avoid believing what is false, and on the other we try to believe what is true.[11] These goals, however, can pull us in different directions. In the cases we've been considering that motivate pragmatic encroachment, given the risk of being wrong, we put more weight on avoiding error than on believing truly. However, sometimes, given the stakes, we should put less weight on avoiding error. For example, whether or not climate change is real has high stakes. If it's real, then we must act immediately to cut back on carbon emissions and fossil fuel extraction. Doing so, however, would have a large economic cost. So, perhaps when the studies first surfaced in the 1980s we did not have sufficient evidence for climate change. However, given the high human costs to not acting—disastrous hurricanes, monsoons, flooding, droughts, etc.—we should have put less weight on the goal of avoiding error. Given the need to act immediately to stem the disastrous effects of climate change we do not need to be 100% certain. So, this cost and risk to being wrong shifts the threshold for justification. Although in most cases it seems to push the threshold up, in some cases it can also push the threshold down. So, we should be careful not to conflate 'high stakes' with a higher threshold in every case.

3.2.6 The General Framework of Pragmatic Encroachment

The takeaway of the intuitive plausibility of pragmatic encroachment is that our practical stakes can determine how good our evidence must be in order to make our beliefs

[10] As we'll see, Pace (2011) develops his account of moral encroachment by drawing on these Jamesian themes.

[11] Schroeder (2012a)'s discussion of balancing type-1 error and type-2 error can be read as a way to attempt balancing these two competing goals.

epistemically justified. Pragmatic encroachment remains, however, controversial. Just as evidentialism faces the challenge of explaining our intuitive judgements in the cases that motivate pragmatic encroachment, pragmatic encroachment shares a fair number of challenges as well. My goal here is not a straightforward defence of pragmatic encroachment. Rather, I present a new challenge for evidentialism and argue that pragmatic encroachment raises a similar challenge to evidentialism. So, pragmatic encroachment can give us a framework that can be extended to capture an important aspect of the way in which we try to hold each other accountable for our beliefs about others.

All of the cases of pragmatic encroachment have focused on trains and banks and how an agent's personal preferences can affect what it is epistemically rational for them to believe. In that regard, all of the cases that have been used to motivate pragmatic encroachment focus on beliefs about objects in the world and only concern the risk of being wrong. However, we do not only have beliefs about objects. We also have beliefs about each other. When it comes to beliefs about others, we need to be cognizant not only of the risk of being wrong, but also of the risk of wronging. This is a challenge that the evidentialist cannot accommodate. I suggest, however, by developing an account of moral encroachment following the same general structure as pragmatic encroachment, we can propose a view that preserves the intuitive plausibility of a weak version of evidentialism while also being able to say that morality places a constraint on what it is epistemically rationally permissible to believe about others.

My goal in this section is to now turn to sketching the most innocuous and least-committal version of pragmatic encroachment I can. Doing so allows us to explore the challenge pragmatic encroachment raises for evidentialism while taking on as few other commitments as possible and thus being neutral between the various theories. Although still controversial, my hope is that by presenting the least-committal version of the basic pragmatic encroachment intuition we can see the shape of the challenge pragmatic encroachment offers against evidentialism, and how it modifies evidentialism in light of the response. And then in the next section we can explore if the same strategy can be extended to accommodate moral considerations. This way, we need not commit ourselves to any particular theory of pragmatic encroachment in presenting a story about moral encroachment.

As we've seen, the pragmatic encroachment challenge to the more restricted accounts of evidentialism tends to begin with a case in which a subject reasonably seems to know, be justified in believing, can rationally rely on their belief, etc. For example, consider the belief that the bank will be open on Saturday. You'd like to deposit your paycheck, but it's not time sensitive that you do so. The pragmatic encroacher then presents a contrasting

case in which the evidential position of the subject is preserved, but the practical stakes have changed. That is, in the contrasting case there is a severe cost to being wrong. For example, it is currently Friday and if you don't deposit your paycheck before Monday your account will be overdrawn. In this contrasting case, you're not sure that the bank will open on Saturday, but the evidence you have that it will be open on Saturday is the same as the evidence you have in the low stakes case in which it is not time sensitive when you deposit your paycheck. Defenders of pragmatic encroachment argue that this cost to being wrong leads us to the intuition that in a low stakes case you can reasonably believe that *p*—e.g., that the bank will be open on Saturday, but that in contrasting high stakes case—despite no change in the evidential position of the subject—it is no longer permissible to believe *p*—that the bank will be open on Saturday. These intuitions clash with *Supervenience* and suggest that *Supervenience* is not as intuitive as initially thought. Recall *Supervenience* says that if the two agents have the same evidence, then if one is justified, so too is the other. In this case, however, our intuitions suggest that the two agents with the same evidence can differ with respect to whether they are justified.

Note, however, that evidentialism cashed out merely in terms of *Support* is compatible with the view that the threshold *t* is determined, in part, by practical interests. If you have a strong interest in the truth of *p*—e.g., there is a large cost to being wrong about p—then the threshold your evidence must either meet or surpass is higher than when that risk of being wrong about *p* is absent. As we've seen, in general there is a question of what determines how much evidence is enough evidence, and all the work in this section has been to show the simple point that practical interests play a role in settling that question. We can now begin to see the shape of an answer to the dilemma posed by Agnes and Spencer; there's some feature of the scenario—the moral stakes—that changes what she and he are epistemically justified in believing. To make this answer plausible, we must first make plausible an extension from pragmatic encroachment to moral encroachment.

3.3 Extending Pragmatic Encroachment to Moral Encroachment

I am not the first to argue that extension can be made. For example, Fritz (2018) argues that insofar as there are good arguments for pragmatic encroachment, there are also good arguments for moral encroachment. Fritz (2018, pp. 657-8) notes,

> Our beliefs are often insufficiently warranted to allow us to rationally take certain high-stakes gambles. But our beliefs are also often insufficiently warranted for us to permissibly act on them when something very morally important hangs in the balance. To the extent that the former phenomenon

> illustrates pragmatic encroachment on knowledge, the latter phenomenon illustrates moral encroachment on knowledge.

However, there are two issues we'd have to confront if we went the same route as Fritz. First, the examples Fritz uses are direct analogues to the traditional pragmatic encroachment cases found in Fantl and McGrath et al., and to rest a defense of moral encroachment on the plausibility of any particular account of pragmatic encroachment risks the plausibility of moral encroachment being undercut by less plausible versions of pragmatic encroachment. Further, such accounts tend to be motivated by a connection between knowledge and action (or justification and action). Notice, however, that I have been careful to motivate the intuitions suggestive of an account of moral encroachment from general observations about the nature of evidence and justification. After all, we've been noting that *beliefs* can wrong independent of their connection to action.[12]

Pace (2011), on the other hand, offers an argument for moral encroachment on epistemic justification that attempts to systematize the Jamesian dictums that we must balance believing truth and shunning error. Following James (1896), Pace suggests that there is no purely epistemic rationale for setting the threshold for justification in one place rather than another. Rather, the evidential standards one adopts, i.e., how one determines the threshold, is determined by and vary according to pragmatic factors. Similarly, although Paul and Morton (2018) do not call their account an account of moral encroachment, they do note that although evidential policies govern the way we adjust our evidential thresholds in different contexts there is no *uniquely* best evidential policy to have. There are multiple evidential policies that are rationally permissible for a given thinker to have from the point of view of purely epistemic considerations. Further, if that's the case, then practical and ethical considerations can and should play a role in deciding between epistemically permissible policies. Consider, for example, their primary case: grit. A capacity for grit can be advantageous for an agent. Further, that it would be advantageous to be resilient to incoming evidence that one will likely fail can bear on the question of which evidential policies one should reason with. A gritty agent requires more evidence before she believes she'll fail than a non-gritty agent; that is, the evidential threshold for

[12]To expand, in these standard accounts of pragmatic encroachment practical considerations encroach *because of* an intuitive connection between knowledge and/or justification and action. The case that I have been using to motivate moral encroachment, however, does not concern action. The case concerns our beliefs about others. This raises worries about whether the account I will be proposing will be a straightforward application of pragmatic encroachment, and these are worries that I am sensitive to. Moral stakes, we will see in Section 3.4.2, encroach in a manner dissimilar to practical stakes. Nonetheless, this remains a stakes-based encroachment account. It just so happens that these observations are also part of the motivation for accounts of pragmatic encroachment.

the question "will I succeed at ϕ-ing if I continue to try?Ó is higher.[13]

Further, in a manner similar to how I will proceed, Schroeder (2018b) argues that that moral encroachment is just a special case of pragmatic encroachment. Moral considerations raise the stakes for how much evidence is required to epistemically justify one's belief. These stakes-related reasons, Schroeder argues, arise from the cost of error. Whereas traditional accounts of pragmatic encroachment have focussed on practical costs, an insight of accounts of moral encroachment is that moral costs of error are a special case of costs that aren't tied directly to *action*. In the standard diet of cases offered for pragmatic encroachment, the costs of error that count against belief are the consequences of what you will do if you rely on the belief. As we've seen, if you have a severe peanut allergy and rely on the belief that your sandwich is almond butter on the basis of insufficient evidence, you risk a terrible turn of events. In these cases, the primary wrong is one of action. However, there can be some moral costs that beliefs carry in and of themselves, i.e., independently of their actual or risked consequences. The argument for the way moral considerations raise the stakes stems from a familiar theme from Chapter 2: interpersonal relationships are constituted, in part, by what we believe of each other. The effects our beliefs can have on our relationships is not mediated by the effects of our beliefs on our actions. The difference, however, between the account I develop in this chapter and the account put forward by Schroeder concerns whether beliefs wrong even when the belief is true. I will say yes, Schroeder no. I will argue that what matters is the agent's belief forming practices and whether those practices, and the basis for their belief, is matched appropriately to the stakes of the situation. Sometimes there can be a morally problematic mismatch when the belief is true, and that the belief is true does not let the believer morally off the hook.

Finally, another account of moral encroachment is offered by Moss (2018, MS). Moss argues that when we use statistical reasoning about a group to conclude something about an individual member of that group, our failure to keep in mind the salient possibility that the individual may be an exception to the statistical generalization is both a moral and an epistemic failing. However, Moss's account concerns beliefs formed on the basis of statistical generalizations and we might worry that there is something objectionable about the statistical nature of the evidence on which the belief in question is based. One might argue that what makes the moral stakes high is itself the statistical nature of the evidence.[14]

[13]Similarly, in defense of epistemic partiality as epistemically rational, Preston-Roedder (2018) argues that because people are *opaque* and the evidence we have about them is at best partial and ambiguous, there are many epistemically permissible responses to evidence. Thus, we are permitted to form judgements about others that cast them in a more or less favorable light.

[14]See, for example, Munton (MS) for a critique of this kind.

Although this is another way to present an account of moral encroachment, if we go this route, we face two questions. First, the question of why statistical generalizations are uniquely subject to moral considerations, whereas forming beliefs on other kinds of evidence isn't.[15] Second, we also face the question of why *some* beliefs formed on the basis of statistical evidence seem fine whereas others don't, i.e., the contrast between profiling pit bulls on the basis of statistical evidence as opposed to people.[16] For these reasons, the cases I'm concerned with extend beyond mere statistical generalizations. Instead, I focus on beliefs that are formed without attention to the moral demands of one's environment and spelling out what that means.

Moving forward I will suggest the following intuitive grounding for moral encroachment. I do not expect all readers to share this intuition, but I do expect that most will recognize the demand, in colloquial terms, to *stay woke*. To be *woke* is to be aware of the moral demands of one's environment. With regard to our epistemic practices, it is the demand to be aware of the moral stakes of our beliefs about one another. It is the demand to be aware of the background against which our epistemic practices exist, i.e., the unjust world we inhabit, and to ensure that our epistemic practices are not only responsive to unjust features of our environment but that they also do not themselves contribute to those unjust features of our environment. Motivated by these considerations, I argue that the epistemic justification of our beliefs is determined, in part, by the *moral stakes*. An upshot of this account is that the epistemic considerations one weighs when trying to settle the question of what to believe include moral considerations.

3.4 The Moral Stakes of Our Beliefs

To get moral encroachment off the ground, one might be tempted by the following strategy. Begin with a case in which a subject reasonably seems to know, be justified in believing, can rationally rely on their belief, etc. In short, a case with low moral stakes. Then contrast that low stakes case with a morally high stakes case in which the subject

[15]Within legal scholarship this is known as the problem of naked statistical evidence. See, for example, Schauer (2006) and Buchak (2014). The common example goes as follows: we know that a bus has hit an individual, but the individual cannot identify the bus. 80% of the buses in the city are blue buses operated by the Blue Bus Company, whereas the other 20% are red and operated by the Red Bus Company. Given the balance of the probabilities, the individual was probably hit by a blue bus. However, this merely statistical evidence is not permissible in the courtroom. Further, if we had an eyewitness who was only reliable 80% of the time, that testimony would be allowed. Thus, we face a problem of explaining why the testimony is allowed, but the statistical evidence is not allowed. I set aside the issues as they appear in the legal scholarship, as my focus is not whether there is something objectionable about beliefs formed on the basis of statistical evidence. Rather, as I've stated, my focus is on beliefs that are formed without attention to the moral demands of one's environment.

[16]To answer this, I do think Moss will need the moral account from Chapter 2.

has a belief, but they are failing to respond with sufficient sensitivity to some morally relevant feature of the environment. Similar to the low stakes case, the belief is supported by the evidence to the same degree. To explain the difference one is forced to conclude that the moral features of the agent's environment are responsible for the change in our intuitive assessment of whether the belief is reasonable or justified. This is despite the fact that both subjects in both cases stand in similar evidential positions with regard to the strength of their evidence. Thus, you get a familiar clash with *Supervenience*. The moral features of one's environment, according to *Supervenience*, shouldn't make a difference to whether it is reasonable to believe. But, just as pragmatic encroachment was compatible with the general form of evidentialism consisting only of *Support*, moral encroachment is similarly compatible. After all, when we are trying to determine how much evidence is enough evidence with respect to a morally weighty matter such as whether you would wrong someone through your belief, such a moral consideration can raise the threshold for the strength of your evidence. Things are not, however, going to be so simple.

3.4.1 Low-Stakes variants

To see this, let us consider *Social Club* again. Plausibly, *Social Club* is a high stakes case. To build a contrasting case, then, we must ask three questions. First, what would the low-stakes variants of the traditional pragmatic cases be; second, what is it that makes these cases high stakes; and third, is what makes these cases high stakes the same in both the pragmatic and the moral domains. To answer the first question, we need a low-stakes case in which an agent is permitted to believe p and a high-stakes case where the agent is not permitted to believe p. The difficulty of building minimally contrasting low-stakes variants for moral considerations will in turn also help us answer the second and third questions. That is enough in the way of suggestive remarks, let us turn to considering the following attempt to construct a low-stakes case and a contrasting high-stakes case.

> *The Enlightened Northerner*. Nathaniel is an enlightened visitor from the North who likes to keep track of how well-staffed social clubs are. He has found that in the South, staff are predominantly black. During his visit to the Cosmos Club, he mentally makes a note of each black person he sees as a staff member, and at the end of day, he leaves for the rest of his journey with the belief that there was one staff member for every ten guests. One of those people he counts as a staff member is John Hope Franklin, the club's first black club member.
>
> *The Racist Club Member*. Duke is a racist member of the Cosmos Club. Ser-

> vice has been bad today, so he wants to reprimand a staff member in front of a lot of people in order to call attention to his dissatisfaction with the service. Furthermore, he wants the race of the staff member to be a major focal point of his complaint. Duke sees a well-dressed black man standing off to the side, and, being a racist, assumes on the basis of his race that the staff member must be lazy and deserving of punishment. Duke then proceeds to confront the man in an attempt to make an example of him.

Despite the acts (and beliefs) of Duke being much worse than Nathaniel, I argue that it is still wrong of Nathaniel to so easily believe of every black person in the club that they are a staff member. Nathaniel is ignoring how his environment is shaped by the South's racism, which here serves as the relevant moral feature of his environment that he is insensitive to. This additionally makes him ignorant of the way he wrongs others by forming beliefs about them on the basis of facts that are due to racism. It is for this reason that both Duke and Nathaniel would still fail to be justified even if the person they picked out is in fact a staff member. Both fail to show the moral care demanded by the unjust nature of the environment they find themselves in. This is not to say that there is no scenario in which they are justified in believing that the person picked out is a staff member. Only that the threshold for justification is higher for that belief in the environment that is the social club than it is in other environments. The reason for that being the history of racism that has resulted in race being a reliable tracker for staff members.

To reiterate, both Nathaniel and Duke fail to show the moral care demanded by the unjust nature of the environment they find themselves in. Whereas the facts may not themselves be racist, the facts that we often rely upon can be the result of racism and racist institutions and policies. Thus, when forming beliefs on the basis evidence that is a result of our racist history, it seems appropriate to ask for more moral care. But what, precisely, does this moral care amount to? It cannot simply be recognizing that the environment is racist. That doesn't seem to be enough—for example, you might think that I can recognize that fact, but then use that very fact as evidence that the black person I'm seeing in the club is a staff member. So, being sensitive to it isn't merely a matter of knowing that we're in a racist environment. Moral sensitivity amounts to not only a mere recognition of an unjust environment, but also an adjustment of one's threshold for justification to meet the moral stakes of one's environment. This is not a fully satisfactory answer, in part, because it turns on the deeper question of what being a moral person requires of us. Nonetheless, we see the following constraint in operation: given the background against which our epistemic practices exist, i.e., the unjust world we inhabit, our epistemic practices should be responsive to unjust features of our environment.

To illustrate with another analogy, consider the following case. You study African-American history and you are putting together a book on poverty in 1930s America. Your focus is specifically on poverty in black communities, so you are visiting archives and collecting photos. The 1930s were a difficult time, and you have very good reason to think that almost any picture you come across of an African-American family will be a picture of a family living in poverty. However, one of the pictures you include is of an African-American family not living in poverty. Contrast this case with one in which you are a photographer putting together a book on North Dakota ranches and one of the pictures included in the book is mistakenly a picture of a South Dakota ranch. There is a wrong in the former case that is different from the latter. It is not merely a matter of presenting things accurately. The problem in the first case is that there exists a neglected demand for extra care and diligence when believing things of others. This same demand is not present when believing things of ranches. I suggest that what is going to separate a high-stakes case from a low-stakes one is that the wrong that's being done has to do with the added layer of a moral injustice—in this case taking the form of a racial injustice.

In answering the first question—what would the low-stakes variants of the traditional pragmatic cases be—the difficulty we face in creating minimal contrasting low-stakes cases for moral encroachment is due to the very features that makes these cases high stakes cases. This, in turn, leads us to an answer to the second question I set forward at the start of this section—what is it that makes these cases high stakes. Underrepresented groups are more often mistaken for employees because of the color of their skin and the racist institutions that make their skin color such a big factor in their inability to gain access to more prestigious employment opportunities. Being mistaken in this context, namely one in which you've historically been excluded, is a greater harm and wrong than being mistaken in a space where that historical disadvantage is lacking. It is this social and moral fact that we've been neglecting. It is this social and moral fact that makes all the cases so far seem like high-stakes cases. I conjecture that this is because of how deeply our social environments are steeped in and shaped by a history of racism. So, whereas we can toggle high stakes for agents in traditional pragmatic encroachment cases by stipulating that they have mortgage payments due, we cannot simply toggle a history of racism on or off to make simple high/low pairs of cases.

To envision such an environment for a low-stakes variant would require invoking something like a twin-earth type scenario where our society hasn't been shaped by the racist events and practices. That instead, this imagined world is such that skin color is as irrelevant to one's social status as [fill in the blank]. Note that '[fill in the blank]' is not an editing error. I challenge the reader to fill in the blank themselves and consider whether

such a society isn't similarly shaped by racist events and practices. For example, suppose we were to fill in the blank with hair colour, height, the ability to roll one's tongue, the length of one's index finger, etc.[17] Plausibly, in such a world any of these options could be relevant to one's social status and the opportunities that one is afforded. Not only am I suspicious of conclusions drawn from highly artificial twin-earth type thought experiments, but also I think to imagine such a world to create a low-stakes variant is to so widen the distance between it and the actual world that the comparison would ultimately have no substantive relevance, thus failing to constitute anything like a "minimal pair".

Bolinger (2017), however, offers an alternative proposal. Consider the following case.

> *Target-Low Stakes*
>
> Agnes and Esther are shopping at Target. Esther wants to purchase the newest iPod for his nephew, and so she needs to find a staff member to unlock the case containing the iPods. They reluctantly wait for a free staff member until Agnes points to a black man dressed in a red shirt and khakis and tells Esther, "There's a staff member that can unlock the iPod case."

Note that the evidence that Agnes relies on in Target-Low Stakes is importantly different from the evidence that Agnes relies on in Social Club in the following regard: wearing khakis and a red shirt is an easily avoidable signal whereas one's skin colour is not. In order to not be mistaken for a staff member in Target, one simply needs to not wear khakis and a red shirt.[18] Bolinger argues that if John Hope Franklin were to wear khakis and a red shirt to Target, his

> voluntary performance of an avoidable behaviour whose social meaning is to signal that [Agnes] is permitted to assume that [John Hope Franklin] effectively waives [his] standing to complain against [Agnes] accepting that [he is a staff member].

Similarly, if I were to go to the Cosmos Club dressed in an attendant's uniform, if someone were to see me and assume I was a staff member, Bolinger argues that the epistemic and

[17] Perhaps it would be wrong to call these 'racist' events and and practices in the imagined world. Perhaps they are instead, hair colorist, heightist, etc. practices. But, that is beside the point.

[18] Of course, it is not that simple. For example, Michelle Obama retells in a People magazine interview,

> I tell this story—I mean, even as the first lady—during that wonderfully publicized trip I took to Target, not highly disguised, the only person who came up to me in the store was a woman who asked me to help her take something off a shelf. Because she didn't see me as the first lady, she saw me as someone who could help her. Those kinds of things happen in life. So it isn't anything new.

moral wrongs of the motivating cases would be absent. Although Agnes's beliefs are still based on a generalization, they're based on generalizations regarding permissible signals (e.g. people who wear staff uniforms) as opposed to morally impermissible signals (e.g. one's skin colour). What is crucial to the alternative diagnosis is that *wearing a uniform* is a public and avoidable signal that one is a staff member. One's membership in a racial group, however, is not a permissible signal. So, in such a case as the following, any moral or epistemic complaint is waived.

> *Lark*. For a lark, John Hope Franklin attends the Cosmos Club dressed in an attendant's uniform. Seeing him, a woman assumes he is a staff member and, handing him her coat check, demands her coat.

Although I would be happy to accept the Bolinger point regarding public and avoidable signals affecting the *degree* of the wrong, it is important to recognize that the historical and social features of an environment *in addition to* whether the signal is public and avoidable play a role in determining whether someone is wronged by the belief.

Further, although in *Lark* there may be less of a moral wrong than in the other *Social Club* cases, there is still some moral wrong here. We cannot ignore the social environment of the Cosmos Club. For example, imagine instead that it is Esther that puts on a staff member's uniform in order to play a prank on Agnes. If she succeeds in being so mistaken, she will not be wronged to the same degree that John Hope Franklin would be wronged if he partook in the same prank. We have good reason to suspect that *even if she were wearing a staff uniform and John Hope Franklin were not*, in the context of the Cosmos Club someone is more likely to ask John Hope Franklin for their coat than they are to ask Agnes.[19]

Bolinger's proposal, then, can help us determine the degree of the wrongs, but it must still be supplemented with a story about the morally relevant features of the environment and how those features affect the orientation we ought to take towards the evidence when forming beliefs about groups that have suffered historical and ongoing injustices in those environments. Supplementing the story in this way also helps with the story regarding which signals are morally permissible. For example, although it is tempting to think that clothing is an avoidable signal so you can make predictions about someone on the basis of their clothing, there are worries about the racialized nature of certain articles of clothing.

As I have argued previously in Chapter 2, we owe everyone the right to not see them as objects to be predicted on the basis of statistical information about them, whether it's their clothing, skin colour, sexual orientation, etc. However, in certain practical environments

[19]Recall Michelle Obama's story of being mistaken for an employee in Target.

we might owe more than in others. That is why in Target we owe less than we owe in Cosmos Club. It is the situation, not the avoidability or public nature of the signal, that triggers the duty and the strength of the duty. What we've seen is that there are many different ways in which our beliefs can wrong. Some beliefs wrong in virtue of not taking the moral standpoint, some beliefs wrong in virtue of ignoring the risk of being wrong, some beliefs wrong in virtue of being based on non-avoidable signals.

3.4.2 What are moral stakes?

One might wonder at this point whether these moral stakes are simply practical stakes. If so, it might seem that moral encroachment is a mere extension of pragmatic encroachment, rather than a view that brings with it anything novel, surprising, or distinctive of its own. This brings us to the third question that I set out at the start of this section: is what makes these cases high-stakes the same in both the pragmatic and the moral domains? Ultimately I will reject the idea that moral encroachment is a mere extension of pragmatic encroachment. However, for now, it will be helpful to explore this possibility. One way to conceive of moral encroachment as a version of pragmatic encroachment is to consider it an extension of practical stakes to others rather than ourselves. For example, consider the following first-pass distinctions between the two versions of the view:

> *Practical Stakes*. A risk of being wrong you impose on yourself when you believe on the basis of insufficient evidence.

> *Moral Stakes (First Pass)*. A risk of being wrong you impose on others when you believe on the basis of insufficient evidence.

Practical stakes concern the agent's preferences, whereas moral stakes take into account other people's preferences. This is a point that has already been recognized. For example, consider the following pair of cases from Ross and Schroeder (2012, pp. 261).

> *Almond Butter-Low*. Five minutes ago, Hannah made three sandwiches and placed them in the refrigerator. She told Sarah that she placed the peanut butter sandwich on the left, the tuna sandwich in the middle, and the almond butter sandwich on the right. Hannah then departed just as Sarah's friend Almira arrived for lunch. Sarah knows that Almira has no allergies. Almira says: "I'd love an almond butter sandwich." And so Sarah opens the refrigerator door, points to the sandwich on the right, and says: "The sandwich on the right is an almond butter sandwich. You can have it."

Almond Butter-High. This case is just like Low, except here it is Sarah's nephew Algernon who is visiting for lunch, and he has a severe peanut allergy. He asks Sarah for a sandwich. Sarah knows that the peanut butter sandwich would be fatal to Algernon, but that the almond butter sandwich would be harmless. She also knows that he would slightly prefer the almond butter sandwich to the tuna sandwich. When Sarah goes to the fridge, she can tell by visual inspection which is the tuna sandwich, but she cannot tell by visual inspection which is the peanut butter sandwich and which is the almond butter sandwich. So she gives him the tuna sandwich.

The contrasting cases above serve as illustrations of how it is not only your own stakes that can change whether you are justified in believing *p*, but also whether someone else is in a high stakes situation with regard to *p* can change your justification as well.

In both cases, Sarah has the same evidence and the evidence supports the proposition that *The sandwich on the right is an almond butter sandwich* to the same degree. If we were still fond of more restrictive accounts of evidentialism we might be tempted to say that if Sarah is justified in believing that the sandwich on the right is an almond butter sandwich in *Almond Butter-Low*, she is also justified in believing that the sandwich on the right is an almond butter sandwich in *Almond Butter-High* (as we've seen, this follows from *Supervenience*). However, given the risk involved with acting on that belief not for herself but for Algernon, our intuitive response is that Sarah is not justified in believing that the sandwich on the right is an almond butter sandwich. So, despite being in the same evidential situation Sarah is justified in *Almond Butter-Low* and not justified in *Almond Butter-High*. As a result, it appears that not only do the stakes for the believer make a difference to whether she is justified, but the stakes for those who would be affected by her acting on the belief make a difference to whether she is justified. But if this is all that moral stakes are, this is uninteresting.

Instead, as I have been suggesting, what is morally at stake in the cases we began with is an added insensitivity to background features of one's environment. This involves a broadening of our understanding of how our epistemic practices are environment-sensitive. There is a morally and epistemically problematic form of insensitivity that I have outlined earlier. This sensitivity is distinctive to our moral agency and it is what separates cases of high practical stakes concerning other people from high *moral* stakes concerning others. Thus, it's critical that we capture this potential for insensitivity in our gloss of moral stakes. Here's a second-pass attempt to capture this point:

Moral Stakes (Second Pass). The risk one imposes on another person by failing

> to respond with sufficient sensitivity to the morally relevant features of the environment when settling a belief.

This gloss on moral stakes focuses attention on what the non-evidential features relevant to settling belief in these cases are. Namely, the relevant features are moral facts that you should be sensitive to. Saying this, however, is pretty empty. It is like answering the question "when do we behave immorally?" by saying "we behave immorally when we fail to respond with sufficient sensitivity to the morally relevant features of the environment before acting." This is not a helpful answer. Obviously what one *wants* to know are what the morally relevant features are that we should be responding to. This is where the work of Chapter 2 will help us. The work of Chapter 2 was an attempt to articulate the morally relevant features we should be sensitive to.

Now, although this bears some similarity to the idea of practical stakes, it cannot be subsumed under our understanding of practical stakes. Recognizing that our environment can place moral demands on our epistemic practices results in a more robust form of pragmatism than is commonly defended. As a result, we should not think of moral encroachment as merely a straightforward application of pragmatic encroachment; rather, it builds upon the foundations laid by pragmatic encroachment and extends our understanding of the non-evidential considerations that matter for justification.[20]

[20]There is a point to be made here about how the cases that are traditionally used to motivate pragmatic encroachment rely on A-stakes (see Worsnip 2015). A-stakes track what is at stake with respect to whether an agent believes (or relies) on the proposition. On the other hand, there are W-stakes which track what is at stake respect to whether the proposition is true or false. That is, the relevant cost is the cost of the proposition of the potential belief being true versus being false as opposed to the cost of relying on the belief should it be false. Worsnip presents the following challenge for accounts of stakes-based encroachment: only a W-stakes reading preserves intuitions about knowledge attributes, whereas only an A-stakes reading preserve the putative link between knowledge and practical reasoning that motivates many pragmatic encroachment accounts. But as we see in light of the cases of interest to this chapter, not only does it seem that in these cases it is the A-stakes that seem relevant thus complicating Worsnip's claim that W-stakes are the only stakes that genuinely relevant for knowledge, but also moral stakes complicate the relationship between W-Stakes and A-stakes.

Perhaps in addition to W-stakes and A-stakes there are also M-stakes that don't quite operate in the same way as either W-stakes or A-stakes. As noted above, W-stakes are the stakes as to whether the proposition is true, whereas A-stakes are the stakes as to whether the subject ought reliantly believe or withhold from believing *p*. M-stakes are neither quite W-stakes or A-stakes. After all, in *Social Club*, it doesn't make sense to say that the W-stakes are high or low. Partly why this is difficult is because Worsnip only presents A-stakes and W-stakes as they pertain to *you*, not to others. For example, you might think that we perhaps could capture the stakes for others in terms the W-stakes. But, Worsnip's W-stakes only consider how the world turning out a certain way rather than another would be bad *for the believer.* Hence, it is unclear to me what the mapping between A-stakes, W-stakes, and moral stakes would be. A conclusion we can draw, however, is that the matter of stakes and what kind of stakes there are, or more difficult than initially thought. There are not only plausibly two kinds of stakes, but potentially three. Thanks to Jessie Munton for pressing me on this point.

3.5 Three Further Worries

Finally, I turn to three objections that can be raised against this account of moral encroachment. First, that moral encroachment is an incomplete account of the ways in which beliefs can wrong, i.e., of doxastic morality. Second, that moral encroachment has perverse consequences. Third, that moral encroachment is too demanding.

3.5.1 The Incompleteness of Moral Encroachment

As Begby (2018) has argued, encroachment-style views are premised on the idea that we can resolve this conflict between epistemic and moral considerations by showing how moral considerations can affect an agent's evidential policies. However, even if the moral stakes or practical considerations determine the evidential threshold a belief must meet, that is a constraint on *epistemic processing*, not a constraint that operates directly on belief itself. As a result, an agent could adopt the morally required evidential policy, adjust their threshold to meet the moral stakes of their situation appropriately, and upon receiving enough evidence, be justified in forming the problematic prejudiced beliefs. In that case, the encroachment theorist must either say that the morally objectionable belief is no longer morally objectionable or, if it is still morally objectionable, it doesn't matter because the threshold has been surpassed. Further, as Marusic and White (2018) have also noted, there is a redundancy problem here. If the belief is genuinely wrong, then the moral questions should remain on the table after the epistemic questions have been settled. Encroachment-style views, then, are incomplete accounts of doxastic wronging.

I have noted previously, however, that there are many different ways in which our beliefs can wrong. Some beliefs wrong in virtue of not taking the moral standpoint, some beliefs wrong in virtue of ignoring the risk of being wrong, some beliefs wrong in virtue of being based on non-avoidable signals. My goal, as I have repeatedly noted, has not been to give necessary and sufficient conditions of when beliefs wrong or a taxonomy of beliefs that wrong. The goal of this dissertation has been more limited, it has been to engage in a particular dialectic with the supposedly rational racist. The goal has been to show how Spencer, by his own lights, fails to live up to the demands of epistemic rationality. This account of moral encroachment is one that Spencer should be able to accept. Spencer should recognize that when it comes to beliefs about his friends, his evidential threshold for believing negatively of his friends is higher than for strangers. Given that it is rational to have higher thresholds in some contexts depending on what you morally owe to the folks in that context, it is epistemically rational for him to set a higher threshold for beliefs about non-dominantly situated groups. Just as he wrongs his friend when he believes on

insufficient evidence that his friend cut someone off in traffic, so too he wrongs the black diner in his section when he believes that the diner won't tip well.

I do leave it open, however, that this threshold can be met, and when the threshold is met the belief no longer wrongs. But, I do not think it then follows that the morally objectionable belief is no longer morally objectionable. There is the belief formed on the basis of insufficient evidence and the belief formed on the basis of sufficient evidence. One wrongs, the other doesn't. To determine which belief wrongs is a matter of how the agent responds to the moral stakes of their situation. The moral concerns are not redundant, they settle the question of what is sufficient.

3.5.2 Perverse Consequences

Here are two seemingly perverse consequences that this account of moral encroachment might have. First, consider people who grow up in deeply prejudiced social settings with no access to contrary evidence. As Begby (2018) argues, these people should be counted as *victims*. If you had the bad luck of growing up in a severely constrained socio-epistemic environment, then holding you responsible and morally blameworthy for the predictable consequences of your limited epistemic opportunities would be perverse. It would compound your plight.

Second, consider what this account of moral encroachment would recommend when it comes to believing things of any member of non-dominantly situated group. You must walk on epistemic eggshells.[21] However, by only walking on epistemic eggshells around members of non-dominantly situated groups and being more free in your epistemic practices around members of dominantly situated groups you may end up compounding a wrong that's done to them. Relatedly, whereas I have focussed on cases in which we wrong by *believing*, when it comes to members of non-dominantly situated groups it seems we more often wrong them when we refrain from believing or when we fail to believe.

To address the second challenge first let us consider an example: when we fail to believe on the basis of a woman's testimony that she has been sexually harassed. It would be inappropriate to say that it's a case in which the stakes are high and so the threshold for settling belief is also higher. What this suggests, however, is that the moral stakes can not only raise the threshold for evidence, but also lower it. I am skeptical, however, that we'll be able to articulate a general principle to determine in every case whether the threshold goes up or whether the threshold ought go down. The very same class of evidence—fact patterns that are the basis or result of an injustice—and the very same way of relating to others—taking the moral standpoint—can sometimes make the threshold go up and

[21] Thanks to Adrienne Martin for suggesting this phrase.

sometimes down. For example, it is easier (and perhaps ought be) easier to believe well of our friends. The threshold is lower. Similarly, perhaps to address the epistemic eggshells concern, when it comes to some situations we should be more ready to believe things of vulnerable groups than otherwise.

Now, to return to the first version of the perverse consequences worry, let me start by noting that although this is a serious challenge I think it brings to light a promising feature of taking the thesis that beliefs can wrong seriously. To see this, let me start with a biographic note. I did not learn about residential schools until I attended university. In no history textbook, in no class or conversation was the fact that First Nation, Inuit and Métis children were removed from their homes and communities in the name of "assimilation" ever taught to me. The way in which I was taught, the way in which all Canadian children were taught, was (and continues to be) constrained by a socio-epistemic environment that deeply wrongs First Nation, Inuit and Métis peoples (see Mills 2007).

An account of doxastic morality that holds people responsible and morally blameworthy for the predictable consequences of our limited epistemic opportunities might be perverse. But, these wrongs at the level of the believer must be balanced against the fact that we must also do justice to this ongoing wrong that occurs when history is erased or forgotten. The wrong to First Nation, Inuit and Métis children is not just a wrong in how the Canadian government acted. If we do not accept responsibility and blame for our beliefs that are the result of forgetting or miseducation, we continue to contribute to a terrifying epistemological reality for the First Nation, Inuit and Métis peoples: we devalue their lives and experiences (see also Dotson 2018).

3.5.3 Too Demanding

Finally, let us turn to the demandingness worry. If one will fail to be justified in virtue of failing to appreciate the burden and risks that they impose on another, then almost all beliefs about other people—especially any belief about a person on the basis of their race or another protected class—are going to be high stakes. This seems overly demanding because it requires moral agents to be fairly sophisticated in recognizing when they should occupy this kind of moral standpoint. As a result, if this level of sophistication is beyond the reach of most, then most (maybe all) of our beliefs about individuals from non-dominantly situated groups are unjustified. After all, it is not clear whether most people would ever find themselves in a position to meet such a requirement. I will address these concerns about sophistication and demandingness by first motivating that it makes sense why such scenarios ought be considered high-stakes and next say a bit addressing why this isn't problematic. To do so, we will find a friend in one of the classic evidentialists: Clifford

(1877).

Clifford is a standard example of an evidentialist, but he is an evidentialist who recognizes the very same puzzling interaction between beliefs and stakes that defenders of pragmatic encroachment attempt to capture. Clifford argues that the rightness or wrongness of a belief turns on not solely whether it is true or false, but rather whether the agent has a right to believe on the evidence before her. To motivate his thesis he gives us the example of a shipowner who casts aside doubts about the seaworthiness of his ship. The shipowner convinces himself that because the ship has returned home safely so many times, he need not worry about whether it would return safely home from this trip as well. But, supposing that the ship does go down, what do we say of him? Clifford suggests that the shipowner is guilty of the death of people on his ship. Further, even if he sincerely and with conviction believed that the ship was seaworthy, "he had no right to believe on such evidence as was before him." Given the high stakes—the possible deaths of the passengers should the ship not be seaworthy—he was not justified in believing that the ship was seaworthy. Further, Clifford argues that even if the ship were seaworthy and the ship had made its voyage safely, that would not diminish the guilt of the shipowner "one jot."[22]

Clifford understands belief as occupying a special social role. It is never the case that any belief is merely trivial or insignificant.[23] He famously remarks that it is wrong always, everywhere, and for anyone, to believe anything upon insufficient evidence. This is a duty we have to all of humanity. Further, if you believe irresponsibly, i.e., without sufficient evidence, then you have sinned against humanity. This moral concern about the social nature of belief leads to an epistemic constraint on how beliefs ought to be formed. It is not just important to the goal of believing the truth that we believe on the basis of sufficient evidence, but also it is our moral duty to do so.[24]

Clifford's view, given how extreme it is, has become an object of ridicule. Famously, James (1896) referred to Clifford as *that delicious enfant terrible*. James accuses Clifford for taking "Shun error!" to be of utmost importance because Clifford is afraid of looking like a dupe if he believes a falsehood. Perhaps James is right, but that is a question for James-Clifford scholars. For the purposes of this chapter, we can latch onto the following thought that we also find echoed by Clifford: there are moral costs associated with believing that should drive us to be more cautious. Given the picture we've been working with, we can see how the moral demands we place on one another extend beyond sim-

[22] We can see Clifford as one of the early proponents of *Uniqueness* with this idea.

[23] Clifford is rightly mocked for this extreme feature of his view. After all, surely the belief that there are an odd number of leaves on the tree outside my window is a trivial or insignificant belief.

[24] See Feldman (2004) and Mitova (2008) for more on this point.

ply how we expect others to treat us, but also what others believe of us. All that moral encroachment requires is a certain level of epistemic humility when it comes to beliefs about others.

At this point, you might be willing to agree that this demanding requirement is unwarranted, but still worry that it requires moral agents to be too sophisticated. Thus, it is demanding in an additionally problematic way: the standard is too difficult to meet. However, given that moral encroachment brings moral considerations to bear on our epistemic practices, this should not be surprising. Morality is demanding. Recall that if we understand moral encroachment as a systematic treatment of the imperative to stay woke, staying woke *is demanding*. One does not simply wake up and stay woke. Rather, it is something that one has to work at everyday. Morality is much the same.

It should not be surprising then that a moral constraint on our epistemic practices would be similarly demanding. In our everyday lives and our day-to-day beliefs we may often fall short of the moral and epistemic ideal. The ideal, however, exists as a standard we ought to strive to meet nonetheless. Being a good person is a moral project that is never complete. Hence an upshot of this account of moral encroachment is that the question of what individuals should believe is, in part, a moral one. The promise of moral encroachment is a promise of reconciliation. Our moral projects and our epistemic projects do not stand at odds with each other. Sensitivity to the moral demands of one's situation not only makes an agent a better moral agent, but a better believer as well.

CHAPTER IV

Assessing Alternatives

Our stalking horse has been the supposedly rational racist, Spencer. Spencer wants to know what's wrong with believing that the black diner in his section won't tip well. The troubling question raised by Spencer, the question that has been the motivation of all the work in this dissertation so far, is the following. If seemingly racist beliefs just reflect reality, what's the way in which they could be either morally or epistemically wrong? The task of Chapter 1 and Chapter 2 was to answer the first disjunct. That is, the task was to show that racist beliefs—even when seemingly supported by the evidence—could *morally wrong*. Spencer, however, could concede this point and note that although he morally shouldn't believe, his belief is nonetheless epistemically rational. Thus, there is still a sense in which he *should* believe the morally objectionable belief. Further, as we've been presenting Spencer, what Spencer claims to care most about is being epistemically rational. To answer this, the task of Chapter 3 was to show that Spencer's belief is not even okay in the respect he cares most about, i.e., the belief is not epistemically justified. As I argued in the previous chapter, if there is good reason to think that non-evidential considerations can affect the threshold we set for when a belief counts as justified, it is also plausible to think that moral considerations are amongst those non-evidential considerations. According to this account of *moral encroachment*, in ignoring the moral demands of his situation, i.e., the relevant non-evidential considerations that affect where we place the threshold for when a belief counts as justified, Spencer fails to see that there is no answer to the question of what makes his evidence sufficient to believe on which he has sufficient evidence to believe. Thus, his belief is both morally objectionable and epistemically irrational.

Moral encroachment, however, is not the only way to answer the challenge raised by Spencer. The task of defending moral encroachment is incomplete until I show why it was worth developing *that* account rather than any other plausible account of what to do in cases of normative conflicts—i.e., conflicts between oughts. The tasks of this chapter,

then, are primarily organizational and clarificatory. With regard to the organizational task, I will present a number of alternative paths we could have taken when presented with what appear to be apparent conflicts between the moral ought and the epistemic ought. I will assess each alternatives' strengths and weaknesses, and consider the answer it offers to the problem of the supposedly rational racist. Doing so also fulfills the clarificatory task of this chapter. In presenting, assessing, and comparing the alternatives with the answer I give in Chapter 3, my goal is to make clearer the motivations and reasons for preferring the moral encroachment view to the alternatives. So, by the end of this chapter we'll have not only a satisfactory answer to the supposedly rational racist, we'll also see why the answer given by moral encroachment—that there is no real conflict between the moral and the epistemic ought when it comes to the question of what one should believe—ought be preferred to the alternatives.

4.1 Conflict Creation and Conflict Resolution Strategies

The two theses and an assumption that set us up with a conflict are the following.

First, according to the evidentialist—whom we discussed at length in the previous chapter—all there is to the rationality of belief, i.e., *epistemic* rationality, is ensuring that you have the attitude that you ought have towards *p*. The attitude you ought have is one that is determined by whether *p* is supported by the evidence. Intuitively, this is a compelling account of what we mean when we say you *ought* believe something. For example, if I report that there's a 92% chance it'll be sunny tomorrow, there is a sense in which you *ought* believe that it will be sunny tomorrow.

A related assumption that the traditional evidentialist smuggles into their account concerns what we mean by 'epistemic'. That is, if we were more careful in the earlier presentation of the evidential thesis, the evidentialist thesis is not simply that the attitude you ought have is one that is determined by whether *p* is supported by the evidence *but rather* that the attitude you *epistemically ought* have is one that is determined by whether *p* is supported by the evidence. This assumption is smuggled in because it doesn't ordinarily need to be defended. If we were to step back and think about epistemology as it is traditionally conceived, i.e., as the study of *knowledge*, then the subject matter of epistemology is the factors in virtue of which a belief is an instance of knowledge.

As Stanley (2005) notes, typically these factors are those that are truth-conducive, i.e., factors that make the belief more likely to be true. If we return now to the question of what you should believe and the rationality of your beliefs, i.e. *epistemic rationality*, what we look for are *truth-conducive* factors to answer the question of what one should believe.

This shouldn't be surprising if we understand epistemic rationality as the kind of rationality that is entailed by knowledge. If it is truth-conducive factors that are relevant to whether a belief is an instance of knowledge, then similarly those same truth-conducive factors are relevant to whether a belief is epistemically rational. As a result, if moral encroachment argues that what makes a true belief into knowledge (or justified) is not entirely a matter of truth-conducive factors but rather a matter of moral and practical factors as well, then on this conception of 'epistemic' it sounds like what the moral encroacher is committed to is rejecting this assumption. That is, what makes true belief into knowledge is *not* entirely an epistemic matter.

But, there is more than one sense in which we can use the term 'epistemic'. As Stanley (2005) notes, there are two senses in which epistemologistÕs use the term 'epistemic'. The first use we saw has to do with factors relevant to whether a true belief is knowledge. Call this 'epistemic$_K$'. Alternatively, 'epistemic' can also be used to denote truth-conducive factors that make the belief more likely to be true. Call this 'epistemic$_T$'. So, when we use the term 'epistemic' we must be careful not to conflate epistemic$_T$ and epistemic$_K$. Recall that I previously noted that the subject matter of epistemology is the factors in virtue of which a belief is an instance of knowledge and that *typically* these factors are those that are truth-conducive. This leaves it open that what makes a true belief into knowledge or makes a true belief rational or justified is not entirely an epistemic$_T$ matter. Returning to the evidentialist, in the previous chapter we saw that the evidentialist defines evidence in support of p as reliable signs, symptoms, or marks of that which it is evidence of, e.g., reasons that raise the probability of p being true. So, we see that the traditional evidentialist restricts epistemic rationality, i.e., the rationality of our beliefs, to being an epistemic$_T$ matter. Moral encroachment, however, is able to recognize a role for both uses of *epistemic*.

The second thesis that sets us up for a conflict is the thesis that beliefs can wrong. If a belief *wrongs*, then there is a sense in which you *ought* not believe. In general you shouldn't violate the demands of morality and the sense in which you ought not believe is this *moral* ought. It follows, then, that if believing that someone won't tip well given the colour of their skin would violate a moral demand, then Spencer, our supposedly rational racist, shouldn't believe that the diner in his section won't tip well. So, we find ourselves facing the following normative conflict: Spencer ought (epistemically) believe p, but Spencer ought (morally) not believe p. When the specter of normative conflict looms, what do we do?

In light of this conflict, in the previous chapter I noted that there are five ways one could go. First, we could explain away the intuitive plausibility of evidentialism insofar

as it seems responsible for generating this kind of conflict. Alternatively, we could show that there is in fact no conflict between the moral and epistemic ought by rethinking and questioning the assumptions that seem to pit them together (as I argued in Chapter 3). Third, perhaps we could show that although there is a conflict between epistemic and moral demands, there is a way of adjudicating between them such that we can say there is something Spencer and Agnes ought all-things-considered believe. Fourth, we could deny that there's anything that Spencer or Agnes ought all-things-considered believe while *also denying* that that means we must resort to the pessimistic conclusion that this conflict is irresolvable. That is, when oughts seem to conflict there is no real conflict, there is one sense in which you should do x and *a different sense* in which you should do y. Finally, fifth, we could just endorse the pessimistic conclusion that there is way to determine what to do in cases of conflicts.

Previously I noted that although the first option might be tempting given the many challenges to evidentialism despite it being considered the orthodoxy, the standard ways of challenging the intuitive plausibility of evidentialism will not help us in this case. My goal in Chapter 3 was the second approach—show that there is in fact no conflict between the moral and epistemic ought by rethinking and questioning the assumptions that seem to pit them together. The task now is to show why it was worth doing the work of the previous chapter by now exploring the alternatives to the encroachment picture I defended. To illustrate this space of alternatives, consider the following flowchart.

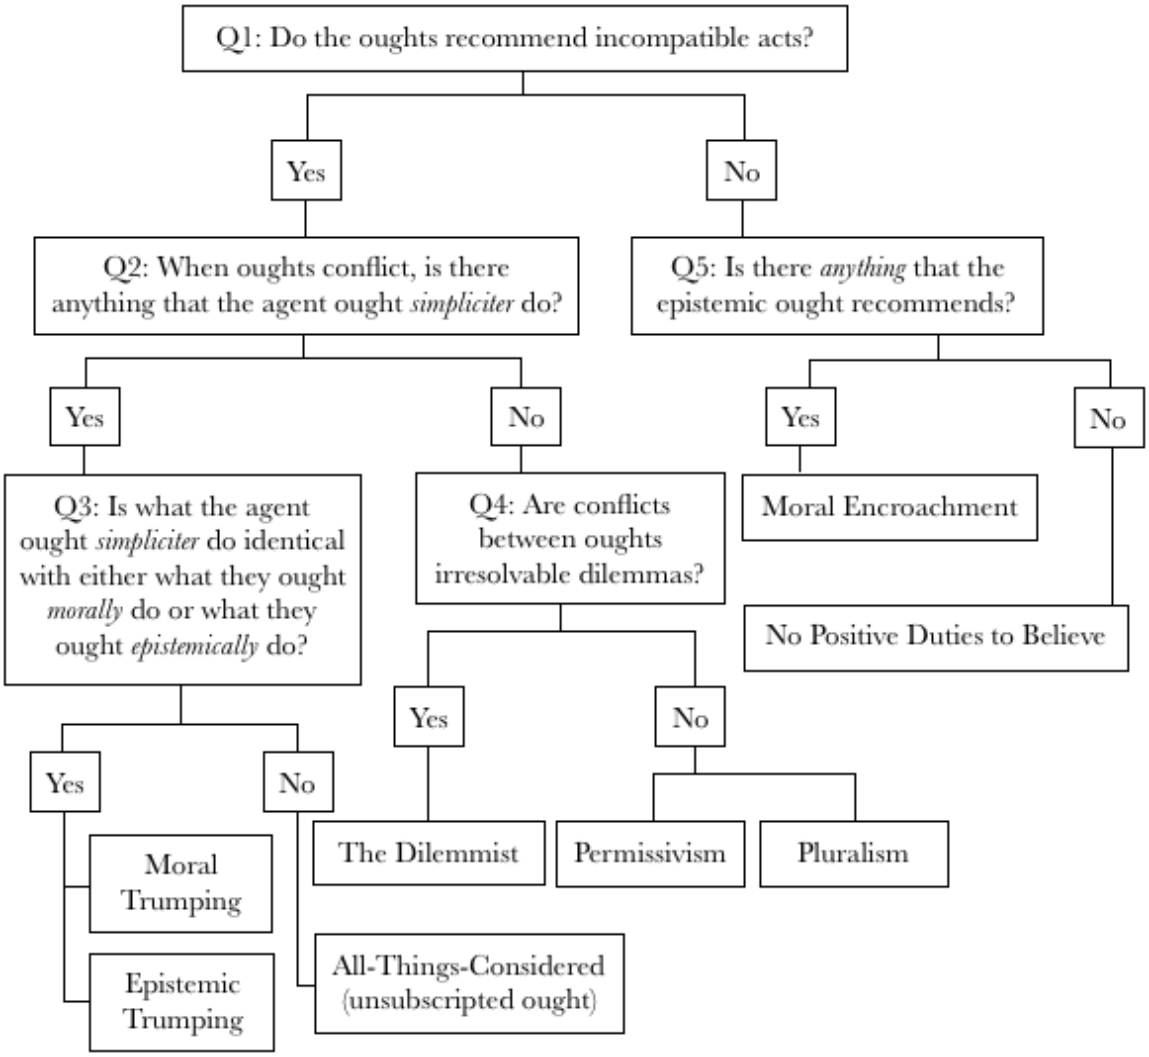

Fig. 2.

As we see in this flowchart, moral encroachment is not even the only plausible account within the space I have been describing under the second option—showing that there is in fact no conflict between the moral and epistemic ought by rethinking and questioning the assumptions that seem to pit them together. That is, there is more than one way to answer "no" to Q1.

Further, as we see in this flowchart, there is a wide range of alternatives to moral encroachment that have thus far gone unconsidered—i.e, accounts that answer "yes" to Q1. These accounts can further be distinguished into two broad categories given how they answer Q2. First, there are the accounts that answer "yes" to Q2, accounts that argue that although there is a conflict between epistemic and moral demands, there is a way of adjudicating between them such that we can say there is something Spencer and Agnes ought *simpliciter* believe. These accounts can further be distinguished by whether

what Spencer and Agnes ought *simpliciter* believe is merely a matter of the moral ought beating out the epistemic (or vice versa) on the one hand, or whether there is an all-things-considered ought that adjudicates such conflicts.

Second, there are accounts that answer "no" to Q2, accounts that deny that there's anything that Spencer or Agnes ought *simpliciter* believe. These accounts can then be distinguished further by how they answer Q4. Either when oughts conflict there is no conflict—that is, there is one sense in which you should do x and *a different sense* in which you should do y—or when oughts conflict there is just no way to determine what to do in cases of conflicts.

I turn now to assessing these alternative accounts, starting with the accounts that answer "yes" to Q1 and Q2.

4.2 The All Things Considered Ought, Ought Simpliciter, or Just Plain Ought

Let us look again at the cases that seem to either set up a normative conflict or are examples of a normative conflicts, i.e., a conflict between different oughts.

> *The Supposedly Rational Racist*. You shouldn't have done it. But you did. You scrolled down to the comments section of an article concerning the state of race relations in America, and you are now reading the comments. The comments on such articles tend to be predictable, but there is one comment that catches your eye. Amongst the slurs, the get-rich-quick schemes, and the threats of physical violence, there is the following comment: "Although it might be 'unpopular' or 'politically incorrect' to say this, I'm tired of constantly being called a racist just because I *rightly* believe of the black diners seated in my section that they will tip worse than white diners in my section." The user posting the comment, Spencer, argues that the facts don't lie, and he helpfully reproduces those facts. For example, he links to studies that show that on average black diners tip substantially less than white diners. The facts, he insists, aren't racist. If you were to deny his claims and were to believe otherwise, it would be *you* who is engaging in wishful thinking. It would be *you* who believes against the evidence. It would be *you*, not Spencer, who is epistemically irrational.
>
> *Mistaken Identity.* The conference has ended, and the organizers have had the forethought to book a number of tables at a nearby restaurant so that con-

versation can continue over dinner. You're having a good time at dinner and, after a few drinks, you get up to use the restroom. As you return to your table, one of the diners, Jim, reaches out to grab your arm and says, "I asked for a refill fifteen minutes ago." For a moment you're confused, then it dawns on both of you what mistake has been made. Most philosophers don't look like you. With regard to melanin levels, you share more in common with the wait staff than your fellow diners. Given your skin colour, the likelihood that you are a member of the staff rather than a fellow diner was high enough to seemingly make it rational for Jim to assume that you were a waiter, not a fellow diner. The belief that Jim had—and in turn his actions—might amount to a social faux pas, but, given that the belief was well-supported by the evidence, the belief—and in turn his act—was reasonable. He's not a bad guy; he just made an honest mistake. Of course, in the moment you don't reason through all of that. In the moment, you quickly laugh it off, eager to return to your table and dinner conversation.

Wounded By Belief. Suppose that Mark has an alcohol problem and has been sober for eight months. Tonight there's a departmental colloquium for a visiting speaker, and throughout the reception, he withstands the temptation to have a drink. But, when he gets home his partner, Maria, smells the wine that the speaker spilled on his sleeve, and Mark can tell from the way Maria looks at him that she thinks he's fallen off the wagon. Although the evidence suggests that Mark has fallen off the wagon, would it be unreasonable for Mark to seek an apology for what Maria believes of him? He is, after all, wounded by what his partner believes of him.

Social Club. Agnes and Esther are members of a swanky D.C. social club with a strict dress code of tuxedos for both male guests and staff members, and dresses for female guests and staff members. They are about to go on their evening walk so they head towards the coat check to collect their coats. As they approach the coat check, Esther looks around for a staff member. All of the club's staff members are black, whereas only a small number of the club members are black. As Agnes looks around she notices a well-dressed black man standing off to the side and tells Esther, "There's a staff member. We can give our coat check ticket to him."

A way to interpret these cases is as though they set up the following kind of conflict:

you ought (epistemically) believe p, but you ought (morally) not believe p.[1] If we take seriously the thought that there really is a deep conflict here, we can see why it would be natural for someone to ask "well, what ought I *really* do?" or "what ought I do *simpliciter*?" I turn now to exploring the alternative accounts that take seriously this idea that there is something that you ought *simpliciter* do. That is, accounts that answer "yes" to Q2: when oughts conflict, is there anything that the agent ought *simpliciter* do?

4.2.1 Ought Simpliciter as Either the Moral Ought or the Epistemic Ought

When faced with such cases, it might be tempting to think that what you ought *simpliciter* do is just what you morally ought do. For a case in which this seems compelling, consider the examples we've seen from Stroud (2006) in which the moral demands of friendship seem to come into conflict with what you epistemically ought believe. Suppose someone tells you a damning story about a friend of yours, e.g., your friend cut them off in traffic. As a good friend, how should you respond to this information? Ordinarily, if you were to receive testimony that some person *a* cut someone else, *b*, off in traffic, that would be sufficient to justify believing that *a* cut off *b*. Stroud (2006, pp. 504), however, argues that we owe our friends "something other than an impartial and disinterested review of the evidence where they are concerned."

In this case of a conflict between the moral ought and the epistemic ought, Stroud suggests that the moral ought *trumps* the epistemic ought. As you consider what you ought *simpliciter* do, you should follow the demands of friendship which require being loyal to your friend and discounting testimony that suggests they behaved badly. As this is an account in which morality wins, let us call it *moral trumping*. If we return to the flowchart, this account not only answers "yes" to Q2, but also "yes" to Q3 (i.e., is what the agent ought *simpliciter* do identical with either what they ought *morally* do or what they ought *epistemically* do? According to *moral trumping* there is something Spencer ought *simpliciter* do, it is just whatever morality recommends. The moral ought is the privileged ought.

On the other hand, another way of responding to these cases of conflicts is to argue that it is not what morality recommends rather it is what epistemic rationality recommends. As this is an account in which epistemic rationality wins, let us call this *epistemic trumping*. The cases in which this alternative seems most compelling take a slightly different form from what we've been considering. Whereas we've been considering cases in

[1] I am careful to hedge my language here because as the reader knows by now, I do not take these cases to be setting up a real conflict between the epistemic ought and the moral ought. The interpretation that there is a conflict assumes that epistemic considerations are completely distinct and nonresponsive to moral considerations, and I have argued previously that this is incorrect.

which the conflict has the form you ought (epistemically) believe but you ought (morally) not believe; the cases that seem most compelling for *epistemic trumping* are cases with the following form: you ought (epistemically) not believe but you ought (morally) believe.

In the literature these appear as cases in which moral reasons to believe seem like *the wrong kind of reason* to believe.[2] For example, suppose an evil demon were to threaten to destroy the world unless you believed that 2+2=5. Morally, you ought believe that 2+2=5, but epistemically you ought not believe that 2+2=5. With regard to these cases, fans of *epistemic trumping* argue that moral reasons are not the *right kind of reasons* for belief. There is something Spencer ought *simpliciter* do, it is just whatever epistemic rationality recommends. The epistemic ought is the privileged ought.

However, with regard to this direction of the conflict—you ought (epistemically) not believe but you ought (morally) believe—there are also cases that seem to make *moral trumping* more compelling than *epistemic trumping*. For example, Preston-Roedder (2013, 2017) argues that civic trust—the kind of trust that is is an essential part of living with others in the sort of harmony that characterizes morally permissible interaction—requires a kind of faith in humanity that is epistemically irrational. Maintaining such an attitude, nonetheless, is an important moral virtue. Further, Marusic (2012, 2015) argues that to make sincere promises we sometimes have to believe against the evidence.

In light of these various cases pulling us in different directions, one might be tempted to posit a separate ought—one that is not identical with either the moral or the epistemic ought, i.e., *ought* as opposed to ought_{moral} or $\text{ought}_{epistemic}$. That is, a separate perspective from which to adjudicate what Spencer ought *simpliciter* do. That is, you might be tempted to answer positively Q1, positively to Q2, but negatively to Q3 (i.e., is what the agent ought *simpliciter* do identical with either what they ought *morally* do or what they ought *epistemically* do?) This brings us a defense of the all-things-considered ought, i.e., a free and unsubscripted sense of ought.[3]

I should note that this gloss on both moral trumping and epistemic trumping is very brief. To fully flesh out both accounts and the pros and the cons for either would require more space. My goal, however, in presenting the range of alternatives that are to the left of moral encroachment (as they appear on the flowchart) is simply to show that there is a range of thorny issues that must be settled and tackled for a full defense if any are to offer a promising alternative to moral encroachment. One key advantage of moral encroachment concerns how moral encroachment can sidestep many of the problems that

[2]See Rabinowicz and Ronnow-Rasmussen (2004); Hieronymi (2005); Schroeder (2012b) et al.

[3]The terminology here is borrowed from Foot (1972), but I should also note Foot as a critic of whether there is a coherent unobscured notion of this kind of ought.

plague accounts that answer "yes" to Q1. As will be shown, the debate to the left of moral encroachment has a troubling circular structure. Troubling objections to the ought simpliciter accounts motivate answering "no" to Q2. However, troubling objections also arise for those accounts, and those objections then lead one back to positing an ought simpliciter. But the previous problems remain which push you back to answering "no" to Q2. Moral encroachment offers a way to side-step these problems. But first, let me continue with the presentation of accounts that posit an ought *simpliciter*

4.2.2 The Free and Unsubscripted Sense of Ought

Let us begin by considering the plain ought that is distinct from the moral trumping and the epistemic trumping views. In general, defenders of the ought simpliciter, the just plain ought, the all things considered ought, whatever you wish to call it, do not identify this ought with the moral ought or the epistemic ought. Partly it is for the reasons we have already seen, i.e., in some cases it seems that it is the moral ought that wins and in others it's the epistemic ought, but sometimes what we want to know is what we *really* ought to do. Another reason is that there are not only moral and epistemic oughts, there are also oughts of etiquette, legal oughts, aesthetic oughts, etc. The all-things-considered ought considers all of these oughts before delivering a judgment. If the ought simpliciter were just identified with one of these oughts, e.g., the moral ought, a consequence would be that anytime we asked about what we ought do we would be asking a moral question. As Wedgwood (MS, pp. 6) notes regarding the all things considered ought.

> This all-things-considered 'ought' is not the same as any distinctively *moral* 'ought'. If there is a distinctively moral 'ought', judgments involving this moral 'ought' would be normative or evaluative judgments of a special kind—judgments based on considerations such as rights and duties, the importance of being fair and helpful towards others, and so on. But not every judgment about what one "ought" to do is a judgment of this kind. For example, if I judge that I ought to buy a new pair of shoes, this need not be a moral judgment. I need not be violating anyone's rights, or neglecting any of my duties, or failing to be sufficiently fair or helpful towards others if I didn't buy a new pair of shoes. Perhaps no one else would be entitled to blame me if I didn't. But it can still be true that I ought buy a new pair of shoes.

In a manner similar to Wedgwood, McLeod (2001, pp. 272-3) notes that when we find ourselves in the difficult situation of trying to adjudicate between conflicts

> ...you might ask—or rather exclaim—"What ought I to do?" That might be a natural question, given the circumstances, but it would be puzzling as well. After all, it seems you already know everything you ought to do: morally you ought to do A, prudentially you ought to do B, legally you ought to do C, and so on. Why then, in the situation described, would you ask what you should do? [...] One reason for thinking this is that, in the situation described at the outset, the question "What ought I to do?" seems reasonable, even though you already know all of your relative duties. Your question would not be reasonable if JPO [just plain ought] were simply identical to one of the relative "oughts." Thus, JPO is not identical to any of them.

Although your moral duty, prudential duty, etc. might, in some cases, also be what you just plain ought to do, McLeod and Wedgwood's point is that this concept of a just plain ought is distinct from these relative oughts.

Now, why should we go in for an account of an ought simpliciter, of a just plain ought? One consideration in favor of such an account is that it offers us a unified answer. Another reason is that the other oughts are partial or incomplete collections of relevant considerations. A just plain ought is comprehensive, it is based on all the considerations that weigh either in favour or against. The moral ought, and other oughts, do this to a degree. For example, when choosing between what you morally ought do and what's in your self-interest, morality isn't blind to your self-interest, it just assigns it less weight.[4] A benefit to the all things considered ought is that it is a kind of comprehensive value (see Chang 2003). Further, this more comprehensive value comes with a kind of special authority that the just plain ought has the other oughts don't. For morality, epistemic rationality, etiquette, prudence, etc. we can always ask "Why should I be moral", "Why should I take my hat off when I enter a church?", "Why should I care about self-interest?", etc. The reason we can ask these questions of these oughts is because we are implicitly granting the authority of a more comprehensive ought, of a third perspective from which to answer these should questions (see Joyce 2001; Velleman 2005).

This view can reasonably be held to be the view of common sense. But as we'll see, there will be many reasons to reject it. But before we turn to those objections, there are still some further points to be made in its favor. As we've seen, there are many uses of 'ought' where ought does not occur in a narrowly moral sense, a narrowly epistemic sense, a narrowly prudential sense, etc. Rather, 'ought' occurs in a more *general* normative sense (see Wedgwood 2007, pp. 24). As Wedgwood (MS, pp. 6-7) notes,

[4]For a dissenting view, see Finlay (2007).

> For this reason, a statement of the form 'A ought to ϕ' need not mean either that A is morally required to ϕ, or that it would best serve A's interests or purposes for A to ϕ; it can mean that A ought to ϕ *all things considered*—i.e., given *all* relevant considerations (which might include both moral and non-moral considerations), A ought to ϕ.

To get an intuitive grasp on this idea, the all things considered ought is the ought that we deploy when we offer advice to one another, it is the kind of ought that would issue from a wise guru who has weighed all the relevant considerations. Consider the following case from (Thomson, 2001, pp. 46) that suggests that there is an unsubscripted all-things-considered sense of ought that could tell Spencer what he ought *simpliciter* to believe.

> Suppose that Alfred is ill, and that only a dose of a certain medicine will cure him. It tastes truly awful, however. Alfred asks us "Ought I really take it?" It is a wildly implausible idea that we can reply only: "Well, your taking it would be very unpleasant, so in one sense of 'ought,' it's not the case that you ought to take it, namely the 'ought$_{enjoyable}$' sense of 'ought.' But your taking it would be good for you, so in another sense of 'ought', you ought to take it, namely the 'ought$_{goodness-for-Alfred}$' sense of 'ought."' It is likely that Alfred will repeat his question: "But ought I take it?" It surely won't do to reply: "Are you deaf? I just told you that in one sense you ought to and in another sense it is not the case that you ought to, and that's all the advice that anyone can give you."

Intuitively, Alfred's question is intelligible. Alfred isn't asking about some relativized ought, he knows what he ought do relativized to each ought. Alfred's also not asking what he should do by his own lights, i.e. what he subjectively ought do, as how would we know better than Alfred as to what he subjectively ought do? These kinds of cases, and those given above, all seem to present a *prima facie* case for an ought *simpliciter* that is distinct from the various relativized oughts such as the moral ought, epistemic ought, etc.

There are, however, numerous reasons to be skeptical of this picture. I will consider some of those reasons in Section 4.3.1, but for now, I turn to considering whether this account should be preferred to that of moral encroachment.

4.2.3 Ought Simpliciter vs Moral Encroachment

According to moral encroachment the oughts do not recommend incompatible acts because both oughts speak to different questions—what the evidence is, and whether the

evidence is sufficient to justify belief—and as such are not at odds with each other when it comes to the question of whether to believe p. It is compatible with this account that there is an all-things-considered ought that also delivers a verdict about what one should believe. It is not clear to me, however, why we must posit such an ought. There is nothing, however, in moral encroachment that commits us to the ought simpliciter view. If you like the ought simpliciter view, however, moral encroachment is compatible with it. I, however, am in general skeptical of the idea that there is an ought simpliciter. Despite the reasons given thus far, I find the idea obscure and mysterious for reasons that I will discuss in the next section.

Another reason in favour of moral encroachment concerns a version of the location challenge that will arise for oughts that are not the moral ought and the epistemic ought. The all-things-considered ought, after all, takes into account self-interested reasons, etiquette reasons, athletic reasons, aesthetic reasons, legal reasons, etc. I spent Chapter 1 arguing that moral reasons are relevant to the question of what one should believe. I argued for this against the objection that moral reasons are irrelevant and that whatever wrong occurs is located somewhere other than in the belief. This objection, which I had answered for moral considerations, can now be raised against the other reasons that the all-things-considered ought takes into account before delivering a verdict about what one ought *simpliciter* believe. That is, reasons stemming from etiquette seem irrelevant to the question of what one should believe. As do athletic reasons, aesthetic reasons, etc. However, we could also imagine cases in which they might bear on what one should believe. If the all-things-considered account says that what Spencer ought believe is determined by taking into account all the reasons for or against belief, it must answer its own version of the location challenge: why do legal oughts, aesthetic oughts, etc. bear on question of what Spencer should believe. Alternatively, why would those considerations have any weight at all?[5]

4.3 Skepticism about Ought Simpliciter

4.3.1 Three Arguments for Skepticism

You may now worry whether there is a coherent notion of this ought *simpliciter* and thus answer negatively to Q2 on the flowchart (when oughts conflict, is there anything that the agent ought *simpliciter* do?). I confess, this is a skepticism that I share. Perhaps there are just facts about what you morally ought do, what you epistemically ought do,

[5]Finlay (2014), critic of the all-things-considered view, also argues that some of these 'oughts' do not have any weight that can be added up. See Chapter 6.

what you politely ought do, what you athletically ought do, what you legally ought do, what you aesthetically ought do, etc. But, there is no ought that adjudicates conflicts between these other oughts. I turn now to considering three reasons for being skeptical about the ought simpliciter. That is, three reasons that push us to answer "no" to Q2.

First, you could worry that these distinctive kinds of reasons or oughts don't share a common justificatory source. For example, Tiffany (2007) argues that there is no common scale on which these various considerations—moral, epistemic, aesthetic, legal, etc.—can be compared. He argues that however we combine these competing reasons is determined by facts about our psychology (see also Finlay 2006, 2009). That is, different agents will weigh considerations differently. There is no further normative perspective from which we can both ask and answer the question of how these reasons should be combined. Tiffany argues that any proposed account of weighings would be guilty of some degree of arbitrariness. On this point, this arbitrariness might be due to the incommensurable nature of these oughts. That is, the oughts are incommensurable, and that is why there is no common scale on which they can be weighed.[6] For example, Kelly (2003, pp. 619) argues:

> In cases in which what it is epistemically rational to believe clearly diverges from what it is practically advantageous to believe, there is simply no genuine question about what one should believe: Although we can ask what one should believe from the epistemic perspective, and we can ask what one should believe from the practical perspective, there is no third question: what one should believe, all things considered. In any case in which epistemic and practical considerations pull in opposite directions, there is simply nothing to be said about what one should believe all things considered.

This is merely a statement of the skepticism, it is not yet an argument for skepticism. To fill in that argument, we can turn to Feldman (2004, pp. 692) who presents the following two examples to argue that there is no sense to be made of a plain ought that "somehow encompasses moral considerations, epistemic considerations, and perhaps others, and then weighs them against one another to come up with an overall assessment." First, consider a child comparing two dolls—where one is short and square, the other tall and thin—how do we determine which one is bigger? Similarly, imagine trying to give an all-things- considered judgement about who is the all-things-considered best billionaire by comparing billionaires according to their wealth and their strength. Feldman argues that it doesn't

[6] I should briefly note that this same worry might push you to posit an all-things-considered ought precisely so that we can weigh these incommensurable oughts. See Chang (2003).

make sense to ask whether height trumps width, or whether volume trumps height, etc. in the case of the dolls. Similarly, there is no independent concept of *strealth* for which there is some correct combination or weighing of the two components—wealth and strength. There is no sense in which one doll is just plain bigger than the other, there is no sense in which one billionaire is just plain strealthier than the other.[7] We could think that this is also the case for the ought *simplicter*. As strealth is the combination of strength and wealth, the ought simpliciter is supposed to be a combination of moral considerations, epistemic considerations and more. So as there is no sense to be made of strealth, perhaps there is also no sense to be made of the ought simpliciter.

Second, you might worry that the ought *simpliciter* doesn't make sense because it requires there to be some standard that is the most normatively important standard and either way you cash out this idea is incoherent. Copp (1997) argues that whatever way we try to explain the authority that this most normatively important standard has will force us to embrace a contradiction. Copp (1997, pp. 101) notes:

> ...the claim that the candidate [standard] S has the property of supremacy is the claim that it is normatively more important than any other standpoint, as assessed from a relevant authoritative standpoint. That is, if S is normatively the most important, then there is some authoritative standard R that yields the verdict that S is normatively the most important standpoint.

The challenge we now face is answering whether R is identical to S. If R is identical to S, then R cannot play the role of the more important normative standard that establishes the normative supremacy of S. That kind of self-endorsement is characteristic of all the normative standards we've been considering and as such it is unimpressive, e.g., morality tells you to listen to morality and self-interest tells you to listen to self-interest. So, R cannot be identical to S.

But, were R to a be a standard *other than* S, that is similarly unimpressive. As Copp (1997, pp. 102) further notes:

> We don't want to know whether S is most important, according to some arbitrary standard R. We want to know if S is most important full stop. But that means R would have to be the normatively most important standard.

[7] I should note that Thomson (2008) has a response to this worry. She argues that normative evaluations are not ambiguous. The best option to choose is which is the most choice worthy option in a particular situation. So in a case like that presented by Feldman (2004) in which are trying to determine who's the best billionaire, that question does not make sense in the abstract. Rather, if we ask for a purpose, e.g., who's the best billionaire to choose for a game of baseball, there is no longer an ambiguity.

> But then S is not the most important standard. This contradicts our initial assumption.

So, Copp concludes, there is no coherent way to cash out the idea that there is a normatively most important standard that is the standard of what we ought simpliciter do.

Third, you might worry, as Baker (MS) does, that if the initial conflict resulted from too many oughts, "how does adding another ought help, instead of multiplying the number of conflicts?" In order for the ought *simpliciter* to add anything and resolve the dilemma it must possess some special property of greater normative authority than the other oughts. Baker (MS) summarizes:

> The characterization of that special property is typically metaphorical or otherwise hopelessly vague and ambiguous. Attempts to eliminate this vagueness face a dilemma: either they characterize the special property in more familiar normative terms, leading to vicious circularity, or they are psychological characterizations that seem to change the subject.

On this point regarding the vagaries of the ought simpliciter, Baker (MS) argues that the ought simpliciter not only lacks any sort of a link to a more familiar theoretical notion—i.e., it is often merely gestured at through the use of evocative phrases—it also lacks theoretical utility. Baker (MS) argues that there is no theoretical problem that the ought simpliciter solves. In the case of the supposedly rational racist, Spencer ought morally refrain from belief, but he also ought epistemically believe. Baker (MS) argues that "[t]here is no mystery here calling out for explanation." Whatever problem there is here is a practical dilemma, and there is no reason to assume that practical problems admit of solution. The problem of the supposedly rational racist, perhaps, is just a practical problem we face given that we live in a world that is structured by racism and racist institutions, and perhaps it doesn't admit of solution.

These reasons for skepticism can lead us in a few directions. We might worry whether denying that there's anything Spencer ought *simpliciter* do amounts to endorsing the pessimistic conclusion that the conflict is irresolvable and he's damned if he believes and damned if he doesn't. In Section 4.3.2 I'll consider the account that answers yes—i.e., that Spencer is damned no matter what he does—and then in Section 4.3.3 I'll consider the accounts that answer no. All of these accounts, I suggest, emerge from a general skepticism of the idea that there is an ought *simpliciter*.

4.3.2 Genuine Dilemmas

The case for genuine dilemmas—a position that, following Thomson (2008),[8] I will call *the Dilemmist*—is most clearly seen in cases of homogeneous dilemmas, dilemmas in which both oughts are of the same normative kind. The primary example of homogeneous dilemmas are moral dilemmas. In this discussion of the dilemmist—who answers the flowchart as follows, Q1: Y, Q2, N, Q4, Y—I will begin by briefly sketching the appeal of the dilemmist position with regard to homogeneous dilemmas and then we can see if it carries over to the case of heterogeneous dilemmas, i.e., dilemmas in which the oughts are not of the same normative kind.

Consider an agent, Ginny, who is torn between two incompatible options. She has promised Fred she'd see the new Marvel movie with him and only him (A), but she also promised George that she'd see the new DC movie with him and only him (B). Looking at her busy schedule she notices that she can only either keep her promise to Fred and do A or she can keep her promise to George and do B. That is, there is no way for her to do both. Ginny finds herself torn between two incompatible options—A and B—and asks "What ought I do? A or B?" Let us imagine that she shouts out this question to the universe, and while she's shouting a dilemmist happens to be walking by. The dilemmist, noting Ginny's dilemma, helpfully answers, "Well, you ought to do A and you ought to B". We could reasonably expect Ginny to reply reminding the dilemmist that she can't do both and she wants to know which she ought to do. If the dilemmist were again to simply reply that she ought do A, and she ought do B. Perhaps she shouldn't do either until she flips a coin.

But, why agree with the dilemmist? Following a line presented by Williams (1965), Thomson (2008) notes that in the case of moral dilemmas you might think that whatever Ginny does, she'll feel regret at not having done the other. If she breaks her promise to George and goes to the movies with Fred, she'll feel the need to make it up to George. Seeing that both incompatible acts are required does justice to why we would feel this regret. If the dilemma could be solved by a flip of a coin, we would fail to do justice to the fact of this regret.[9]

So now let us reconsider the heterogenous dilemma that has been our focus and see whether this fact of regret carries over. As Gendler (2011) introduces one of the cases that has been our focus, *Social Club*, let us reconsider what she says. Gendler suggests

[8] Although I am using Thomson (2008) to present the dilemmist position, Thomson herself rejects the dilemmist position.

[9] Thomson (2008) contrasts this kind of *hard dilemmist* with a *soft dilemmist* who argues that what Ginny ought to do is ambiguous. This would be compatible the pluralist and permissivist accounts will be laid out.

that given that we live in a society structured by racial categories we simply face a tragic irresolvable dilemma. She presents the case of *Social Club* as one in which we face the following dilemma: we must either (a) pay the epistemic cost of failing to encode certain sorts of base-rate or background information about cultural categories (i.e., encoding the information that a minuscule fraction of club members are black whereas all the staff members are black) or (b) we must actively believe against the evidence and work at regulating the inevitable associations which that information gives us. In other words, we must choose between doing what we epistemically ought (encode the base-rate information about the race of the staff members) and what we morally ought (not use someone's race as a determiner for whether they are staff even when it's the determiner that is highly correlated with correctly picking out a staff member). This, she argues, places us in a tragic irresolvable dilemma. Not only can we not do both what morally ought and what we epistemically ought, there is no way to resolve this conflict.

Now, is it the case that as with Ginny's dilemma a simple coin flip could settle for Agnes or Spencer what they ought believe? I do not think so. Imagine Agnes approaching John Hope Franklin as though he were a staff member, being corrected, and then replying "Well I was torn between what I morally ought do and what I epistemically ought do, so I flipped a coin and went with what I epistemically ought do." From Chapter 2, it seems that John Hope Franklin can rightly feel *wronged* by that, and in general, when you've wronged someone you generally feel regret about it. So, this fact of regret carries over to the case of heterogeneous dilemmas. That there is this feeling of regret, this feature than an agent would be wronged if one simply flipped a coin to make a decision concerning whether they are a staff member, suggests that perhaps we a face a genuine dilemma in these cases. If either option was permissible and the dilemma could be settled by flipping a coin, then John Hope Franklin wouldn't be in a position to feel wronged. Notice that we can avoid these grounds for regret by refraining from belief (as suggested by moral encroachment). The dilemmist, however, seems committed to saying that you should both believe and not believe.

As we'll see, however, there is a defusing move that is available if we take a slightly different track on the flowchart. The pluralist, whose account I will expand upon in Section 4.3.3, notes that given the plurality of oughts, these oughts are constantly in conflict. There is nothing special to lose sleep over with regard to conflicts between moral oughts and epistemic oughts any more than we should lose sleep over conflicts between athletic oughts and aesthetic oughts. No standard is more normatively important. As we'll further see in Section 4.3, all of these accounts that emerge from a skepticism concerning ought simpliciter result in this kind of pessimistic conclusion. Spencer, our supposedly

rational racist, will always end up doing something right, even if he also does something else wrong. This is unsatisfying, and I will explain why.

4.3.3 Permissivism and Pluralism

First, this skepticism about an ought *simpliciter* does not require endorsing that these conflicts are irresolvable. That is, if we return to the flowchart, it is possible to answer "yes" to Q1 and "no" to both Q2 and Q4. Instead of the dilemmist picture we saw articulated earlier, one could endorse a kind of *permissivism* about the oughts in these cases of normative conflict. It is not the case that Spencer ought *simpliciter* believe that his neighbour has an open arrest warrant (call this option *A*), nor is it the case that Spencer ought *simpliciter* not believe that his neighbour has an open arrest warrant (call this option *B*). Rather, there is just some sense in which he is permitted to do either *A* or *B*. Both incompatible options are permitted (see Horty 2003 and Brink 1994). However, what this account seems to get wrong in the case of the supposedly rational racist is that it is not that Spencer is permitted to believe or not believe, but there's a real sense in which he gets something wrong when he reasons in this way. But, if you don't feel the force of that you likely don't feel the force of the thesis that beliefs can wrong and there is nothing more that I can say to convince you if you are drawn to this permissivist line.

Second, one could deny that both options are permitted, but insist that there are a plurality of oughts and from the perspective of each ought you simply ought do what it prescribes. As Case (2016, pp. 2-3) presents pluralism, pluralism is committed to the following three claims.

1. *Source pluralism:* There exists an irreducible plurality of normative domains, which issue oughts of distinct kinds.

2. *Conflict*: It is possible for two or more of these domains to issue conflicting recommendations on what to do (i.e., "X-ly one ought to ϕ, Y-ly one ought not to ϕ," where X and Y are two normative domains.)

3. *No authoritative adjudication:* There is no "all things considered" domain capable of issuing recommendations about what one ought *simpliciter* to do; hence, no authority exists that is capable of adjudicating disputes between normative domains.

According to pluralism there is an irreducible plurality of normative domains, and although these domains sometimes issue conflicting recommendations, there is nothing that one ought *simpliciter* to do. Things bottom out at what each respective domain says you

ought to do. So, in the case of Spencer, he ought morally not believe and he ought epistemically believe. It is not, however, as though Spencer is damned whatever he decides to do. There is simply nothing more to say. The accounts that I listed under skepticism about ought simpliciter are all pluralist accounts, and I turn now to expanding further on the accounts defended by Tiffany (2007); Copp (1997); Baker (MS); Finlay (2014) et al.

Pluralism comes in two varieties: unrestricted and restricted. As Case (2016) notes in his arguments against pluralism, according to unrestricted pluralism all coherent standards are reason-generating normative domains, whereas according to restricted pluralism only some standards are reason-generating normative domains. For example, suppose you are a member of the Knights Who Say Ni. One of the standards governing Knights Who Say Ni is to shout "Ni!" until your demands are met. For the unrestricted pluralist, this standard is reason-generating for you. So, in the same way that as an agent that cares about morality you have moral reasons not to murder, insofar as you are a Knight Who Says Ni you have reason to shout "Ni!" until your demands are met. According to a standard example of an unrestricted pluralist, Tiffany (2007, pp. 255), any standard can be a reason-generating standard as long as:

> There is some well-defined aim (e.g., intrinsic-desire-satisfaction), institutional framework (e.g., positive law), or standard of value (e.g., aesthetic value) relative to which considerations may be judged as standing in some normative relation to action (e.g., favouring, defeating, enabling)

In contrast restricted pluralists, such as Copp (1997) and Stroud (1998), argue that these Knights-Who-Say-Ni reasons shouldn't be on par with reasons of morality, epistemic rationality, prudence, etc. Stroud (1998) argues that reasons like the Knights-Who-Say-Ni reasons are domain-relative and they have no normative force on their own.[10] These reasons might have instrumental value, e.g., social-climbing reasons are reasons that are instrumental to the goal of social climbing, but they are only D-reasons, i.e. domain-relative reasons. These reasons are not *genuine* reasons for action.

The unrestricted pluralist could reply that the reason it *seems* as though there are not genuine Knights-Who-Say-Ni reasons is because most people aren't partisan toward that domain. Our partisanship might affect what we take to be normative reasons. As Tiffany (2007, pp. 244-5) notes.

> Just as one may be a partisan supporter of the Canucks over the Maple Leafs—perhaps even seeing support for the Leafs as a character flaw, admittedly non-

[10]For other examples, Dorsey (2013) gives the example of reasons to follow the Constitution of the Satanic Grave Robbers Society, and Copp (1997) gives the example of moon-love reasons.

> culpable for those raised in greater Toronto—without thinking that there is some deep metaphysical truth backing up one's partisanship; so too can one be similarly partisan toward, e.g., morality, prudence, or authenticity.

A problem here for this unrestricted version of pluralism is that this view seems to devolve into nihilism about practical reason and a skepticism about normative importance. As Case (2016) notes, for every coherent standard (and there are many according to unrestricted pluralism), there is an equal and opposite standard that is also reason-generating. As a result, we will have an infinite number of reasons generated by an infinite number of standards and we will constantly be pulled in every direction. Case (2016) argues that if all competing standards are rightly thought to be normative domains, then the result is that none are. Case (2016, pp. 6) notes:

> A normative domain's recommendations must deserve our respect, and must possess authority over the recommendations of competing standards at least some of the time. I cannot make sense of the claim that all standards are normative domains any more than I can make sense of the claim that everyone is properly addressed as Thakin, if this is truly a title of distinction. The only sense I can make of the suggestion that all standards are normative domains is that no standards are really normative domains; hence unrestricted pluralism, so understood, collapses into nihilism.

If the pluralist attempts to avoid this collapse into nihilism by arguing that partisanship makes a normative difference, i.e., that the only domains that are normatively significant for me are ones toward which I am partisan, then pluralism devolves into subjectivism about practical reason. Both nihilism and subjectivism about practical reason are incompatible with pluralism. Nihilism about practical reason is inconsistent with *source pluralism*, i.e., that there are a plurality of normative domains (as according to nihilism there are *no* normative domains). Subjectivism is inconsistent with *no authoritative adjudication* because one's partisanship would determine what standards are authoritative for one.

This might lead you to reject unrestricted pluralism for restricted pluralism, but there are issues for restricted pluralism as well. Recall Stroud (1998) argues for domain-relative reasons that contrast with genuine reasons for action. The restricted pluralist now faces the challenge of explaining why some standards are normative and deserve to be respected, whereas others do not. Further, the restricted pluralist must also explain why our intuitions in cases involving a small good in one domain at the expense of a large cost in another domain suggest that there is something you plain ought to do. For example, if you are walking by a lake and see a drowning child you could either save the child or you could

not. From the perspective of self-interest, perhaps you shouldn't save the child—after all, you would get your clothes wet and from the perspective of self-interest you prefer having dry clothes. From the perspective of morality, you should rescue the drowning child. The pluralist in this case must simply say that what you morally ought do is rescue the child, but self-interestedly you ought not rescue the child. This seems like the wrong conclusion. What we want to say in this case is that you just ought to do what you morally ought do. Thus, we are pushed back towards positing an ought *simpliciter*.[11]

This intuition we have is also the intuition we have in the case of the supposedly rational racist. Just as it doesn't seem right to say of the person that didn't save the drowning child for self-interested reasons that at least they did what they self-interestedly ought to have done, it similarly seems wrong to say of Spencer that he at least did what he epistemically$_T$ ought to have done.

Recall the apology that Jim offers to the conference goer he mistakes for a diner from Chapter 2: "My evidence supported that you were a waiter, not a conference participant. So, I had the attitude I epistemically ought to have had towards you being a waiter as supported by my evidence. This belief wasn't motivated by any feeling of ill-will toward you, I was just believing in accordance with the evidence." As I argued, in this apology Jim problematically takes the theoretical perspective towards the agent wronged by his belief. In his apology he is responding as though he is a scientific observer explaining casual phenomenon in the world. But, as I argued, that is precisely what he should be apologizing for. He should be apologizing for failing to take the participant attitude. For failing to see why you would be invested in the attitudes he holds of you. For failing to be sensitive to the moral considerations and seeing how those moral considerations bear on what he should believe. The pluralist offers a similar response to the one offered by Jim. According to the pluralist, Jim is off the hook for any criticism of the form "you shouldn't have believed that of the diner" as long as he does what he epistemically ought. We could criticize him for not doing what he morally ought, but epistemically he's off the hook. The advantage of moral encroachment is that it can explain why Jim isn't off the hook epistemically either.

4.4 No Positive Duties to Believe

Finally, let us return to the first question on the flowchart—do the oughts recommend incompatible acts?—and consider another view that answers "no" to that first question,

[11]This "argument from notable-nominal comparisons" can be found in Chang (1997); Parfit (2011); Dorsey (2013) and Scanlon (1998). Case (2016) also discusses this problem and raises two further objections to restricted normative pluralism: the sorting problem and the concurrent-case argument.

but nonetheless can be distinguished from moral encroachment by how it answers another question. That is, the question of whether there is *anything* that the epistemic ought recommends (Q5). Moral encroachment answers "yes" to Q5 because the epistemic ought concerns the strength of one's evidence. That is, we have epistemic obligations to attend to the evidence. However, attending to the evidence is not sufficient to settle the question of what one ought believe. Nonetheless, there is something the epistemic ought recommends according to moral encroachment, the Cliffordian charge "to attend to the evidence!".

However, another way to proceed is to answer "no" to Q5. That is, to deny that there is a conflict in these cases because the epistemic ought does not recommend or require anything. One such account is provided by Nelson (2010). Nelson argues that we have no positive epistemic duties to believe, that is, there are only negative epistemic duties. Nelson argues that the most authority the epistemic ought has is in saying what we are permitted to believe. What settles whether *to believe* or might require one to believe, are non-epistemic considerations. This, so far, sounds a lot like the moral encroachment picture that I have been defending. As we'll see this picture is compatible with moral moral encroachment, but there will be one important point of difference.

4.4.1 The Account

Nelson argues that there is nothing we ought believe on purely epistemic grounds. His argument starts by contrasting ethics and epistemology. In ethics, he argues, it is hard to deny that there are both positive duties—things we ought to do—and negative duties—things we ought not do. There is more debate, however, in epistemology regarding whether there are things that we ought to believe and things that we ought not believe. A reason for initially thinking that there are positive epistemic duties stems from the many parallels between ethics and epistemology. First, both ethics and epistemology are *normative*. That is, both disciplines concern what we ought do. Further, many of the meta-frameworks, normative theories, and accounts of structure we employ in ethics have counterparts in epistemology as well, e.g., non-cognitivism, reductionism, non-naturalism, nihilism, consequentialism, deontology, virtue theories, foundationalism, coherentism, etc. As a result, we could suppose that "whatever is true in ethics, about action, is also true, mutatis mutandis, in epistemology, about belief" (Nelson, 2010, pp. 84).

In short, the basic thought for there being positive epistemic duties is that there are positive duties in ethics and perhaps those positive duties extend into epistemology as well. However, as Nelson (2010) persuasively argues, the idea that we have *positive* epis-

temic duties seems demanding and unrealistic. Evidence might give us a reason to believe something, but to say that it follows that I have a duty to believe everything for which I have evidence leads to a conclusion we ought reject: that I am required to believe an infinite number of things. This is what Nelson calls "the infinite justificational 'fecundity' of perceptual and propositional evidence."

This, Nelson (2010, pp. 87) argues, is suggestive of a permissive epistemic theory in which "first-order normative epistemic principles concern what we are permitted to believe, given our epistemic circumstances—not what we are obligated to believe." According to this permissive epistemic theory certain beliefs are licensed by our epistemic situation, i.e., the evidence, but what we should believe is constrained by other non-epistemic considerations. Note that this is entirely compatible with moral encroachment according to which evidential considerations set limits on what what we are epistemically permitted to believe—i.e., the strength of the evidence determine what you are permitted to believe—while the non-evidential considerations, such as moral considerations, determine what you should believe.[12]

Before comparing these accounts, however, let us turn to how Nelson establishes his conclusion. For consider the following case.

> Given the appearance of some distinctive dark, winged shapes, moving across my visual field, what should I believe? That visual evidence, joined with other factors, may license me to believe propositions such as:
>
> (1) There are things moving through the air in front of me
>
> (2) There are birds flying in front of me
>
> (3) There are jackdaws flying in front of me
>
> (4) At least three jackdaws exist (Nelson, 2010, pp. 87)

Of these, which proposition I believe will depend on many things. For example, it could depend on "how my perceptual abilities have developed (e.g. have I learned to discriminate different kinds of bird on the wing?); the background information I happen to have (e.g. do I know what a jackdaw is?); and my particular interests at that moment (e.g. what do I want to know or do now?)."

[12]We have also seen similar remarks made by Paul and Morton (2018) where they note that there are multiple evidential policies that are rationally permissible for a given thinker to have from the point of view of purely epistemic considerations. Similarly, Preston-Roedder (2018) argues that because people are *opaque* and the evidence we have about them is at best partial and ambiguous, there are many epistemically permissible responses to evidence.

When we ask which proposition I should believe, the answer depends only partly on my epistemic situation, i.e., the evidence. The answer also depends partly on my needs and interests. For example, if I am interested in launching a model airplane, then I should believe (1). But, if I am afraid of birds and wish to avoid them, then (2). If I am conducting a species survey, then depending on the specificity of the survey I should believe either (3) or (4). However, if my interest is to hail a cab, I need not believe either (1)-(4).[13] Thus, the difference between ethics and epistemology, is that although there are often things you positively ought do given the morally relevant features of your environment, there is nothing you positively ought believe given the totality of epistemically relevant features of your environment.

I should note here that Nelson is using "epistemically relevant features of your environment" in a way that stands at odds with the claims I make regarding moral encroachment. That is, Nelson is using the restricted 'epistemic$_T$' version of 'epistemic'. According to moral encroachment, the epistemically relevant features of your environment include moral features of your environment, not the merely truth-conducive features of your environment. However, if we make this kind of change, then Nelson's view just becomes the moral encroachment view. As I have argued, I do not think there is any good reason to exclude non-evidential considerations from the list of considerations that are epistemically relevant. Once we include those considerations, however, there will be positive epistemic duties to believe. To clarify terminology once more, by positive epistemic duty here, I will mean a reason to believe. According to moral encroachment both moral and epistemic reasons, insofar as both can count as reasons to believe, will be classed as positive epistemic duties. And this again is compatible with Nelson's view as the evidence alone does not settle what you should believe, to have positive duties the non-evidential considerations must enter into the picture.

Returning now to Nelson's picture, what explains this difference between ethics and epistemology according to Nelson? The answer brings us back to the infinite justificational fecundity of evidence. Nelson (2010, pp. 96) notes that the infinite justificational fecundity of evidence is

[13]Further, although this example concerns perceptual beliefs, the general point extends to inferential beliefs as well. Nelson (2010, pp. 87-8) notes:

> what conclusion (epistemically) ought I to draw if I believe 'p' and 'if p then q' (and grasp the relevant logical rules etc.)? It is impossible to say in advance. It may be 'q', of course, but depending on my needs or interests it may equally be 'q or r', or 'p', or 'p and if p then q'—or, if I am looking to hail a taxi, nothing at all. The premisses license all of these and more, they constrain me from believing anything incompatible with the licensed beliefs, but they require me to draw no conclusion on any topic that does not concern me.

> ...the fact that every single bit of evidence, whether experiential or propositional, potentially epistemically justifies an infinitely large array of different beliefs. By itself, this may not appear terribly significant. When combined with positive epistemic duties, however, it takes on a different appearance, because together they entail the duty to believe everything that is justified for us in our epistemic circumstances.

According to this result, the positive epistemic duties thesis—you ought add every propositional belief that the evidence epistemically justifies for you—and the fecundity thesis—the evidence justifies infinitely many propositions—you ought now add infinitely many beliefs.[14] But as Nelson (2010, pp. 98) notes:

> we cannot have a duty to add an infinite number of beliefs at any given moment, because (given similarly reasonable assumptions about our limited psychological capacities) this is not humanly possible, and we do not have a duty to do the impossible. Hence, we have no positive epistemic duties.

It turns out that belief, unlike ethical action, is optional.[15]

4.4.2 The Drawbacks

One might think that an advantage of this Nelson picture over that of moral encroachment is that it keeps the distinction between epistemic considerations and moral considerations. According to moral encroachment, there is no conflict because part of the epistemic considerations—where by 'epistemic' we mean whatever considerations are relevant to whether a belief is epistemically justified—are the moral considerations. The moral considerations play a role in setting the threshold the evidence must pass for a belief to be to justified. That is why Spencer's belief is not epistemically justified in the way that he thinks that it is. On Nelson's picture, however, Spencer is epistemically justified, he's just not required to believe what he thinks he's required to believe, he's merely permitted to believe it. If Spencer were to rest his defence of believing that the diner in his section won't tip well on the grounds that he's not doing anything wrong because he's doing what he's epistemically permitted to do, on Nelson's account we can question Spencer

[14] An earlier version of it can also be found in Feldman (2004, pp. 678): "Suppose that a person has evidence that conclusively establishes some proposition, q. There are then a huge number, perhaps an infinite number, of obvious logical consequences of q that are also supported by this evidence."

[15] I should also note here that there are other ways to avoid Nelson's conclusion here. You could argue that it's not the case that every belief we have must be represented in the mind somewhere. Many beliefs can be implicit, can be forgotten, etc

about why he chooses to believe that of his neighbour, but not believe all the other beliefs he's permitted to believe. What he chooses to believe, then, could be a reflection of a morally criticizable character akin to the story that Arpaly (2004) and co. tell about the supposedly rational racist back in Chapter 1—i.e., Spencer is likely engaging in motivated reasoning. As I previously noted in Chapter 1, a disadvantage of this approach is that we must say that Spencer's self-conception of himself—as someone earnestly just believing in accordance with the evidence—is false. This may in fact be the case for many people who claim to be like Spencer, but there's something more powerful about an account that is able to preserve Spencer's self-conception of himself and still say why what he is doing is not right *even by his own lights.*

A further drawback of Nelson's account is that no matter how strong the evidence you are never required to believe. This, however, seems counterintuitive and seems to license a kind of epistemic behaviour we might call *wishful denial*. Wishful thinking is when an agent believes that something is the case despite evidence to the contrary because believing would make their life go better in some way. For example, I might wishfully believe that I am not as behind on my work as I in fact am because that belief helps me avoid crippling anxiety. On the flipside, wishful denial would be avoiding believing something you ought believe. For example, I ought believe that I am in fact behind on my work, but I can use Nelson's argument to show that I'm not required to believe that I am behind on my work. For such a trivial case, perhaps that is fine. But now consider more pressing cases that we previously considered in Chapter 3, Section 3.5.2. Consider, for example, the tobacco industry insisting for years that contrary to the evidence we had no reason to believe that there's a connection between smoking and cancer. Or business and financial interests insisting that contrary to the evidence we have no reason to believe that climate change is happening. Or when a woman comes forward with an account of sexual harassment saying that we have no duty to believe her. In such cases, it seems compelling to say that someone who fails to believe that there's a link between smoking and cancer, or a link between the increased carbon emissions from burning coal and the sea levels rising, or who refuses to believe testimony concerning sexual harassment is doing something epistemically wrong. In this cases it seems that they *ought* believe.

Nelson might reply that that is because of the moral stakes of the beliefs, so again it is something non-evidential that is responsible for the ought. But I suggest that even if the moral stakes were absent, the sheer strength of the evidence would be sufficient to force an agent to believe. To stubbornly refuse to believe in the face of irrefutable evidence strikes me to be an epistemic fault. The benefit of the stakes picture is twofold. First, in contrast with Nelson, although moral considerations can shift the threshold once that threshold

is crossed you *ought* believe. Second, the moral stakes can not only raise the threshold for evidence, but also lower it.

4.5 Concluding Remarks

My goal in this last chapter has been to complete the task of defending moral encroachment. The goal of Chapter 1 and 2 was to show that beliefs can wrong, and the goal of Chapter 3 and 4 was to show that the beliefs that wrong are also epistemically irrational. Although I developed an account of moral encroachment in Chapter 3, that account was incomplete until I could show why moral encroachment should be preferred to the traditional ways of answering conflicts that arise between different oughts. As I hope to have shown in this chapter, a key advantage of moral encroachment is that it allows us to prescind from questions about whether there is an all-things-considered ought that adjudicates conflicts and thereby avoid many of the challenges that arise for defenders of the all-things-considered ought and those skeptical of it.

CONCLUSION

A central theme in this dissertation has been that when it comes to what we should believe, morality is not voiceless. Our epistemic practices exist in an unjust and non-ideal world and there are moral and social constraints on our epistemic practices. Specifically, I have argued that beliefs can wrong and the epistemic justification of our beliefs can be affected by the moral demands of our environment. There is much, however, that still needs to be said and this conclusion is not so much a final word on the topic, but rather a confession of everything else I wish I could have addressed.

For example, the focus of this dissertation has been restricted to directed wrongs, i.e., wrongs done to specific individuals. This leaves open, however, large questions concerning when the wrongs done to an individual extend to the group the individual belongs to. For example, when you believe of an individual that they are dangerous given their race, that wrong plausibly extends to the group. As Kendrick Lamar raps in Fear, "How they look at me reflect on myself, my family, my city". This may not, however, apply to other groups, e.g., employees and their place of work. This is a difference in need of an explanation and sadly it's not something I had time to take up in this dissertation.

I have argued that we can wrong each other simply in virtue of the beliefs we hold about each other, regardless of how or even whether those beliefs manifest in our actions. To what extent, however, can I legitimately place a claim on you to believe a particular thing about me? A challenge for this account that I have not yet addressed concerns giving an adequate explanation of what we can legitimately require of othersÕ beliefs about us (see also Begby 2018 for this line of objection). For example, I may want you to believe well of my philosophical ability, but what right do I have to claim that of you? Do I really have a moral claim upon you to believe that I am an excellent philosopher? For a more troubling example, consider Srinivasan (2018)'s (2018) discussion of Elliot Rodger. Quoting from his manifesto, she notes that he claims to have been "cast out and rejected, forced to endure an existence of loneliness and insignificance, all because the females of the human species were incapable of seeing the value in [him]." What right does Elliot Rodger have to demand that the females he interacts with see value in him? In particular, notice that where Elliot Rodger speaks of value, what he means is being sexually desired by "hot, beautiful blonde girls". Certainly he has no right to that. In short, when do we make morally permissible demands on the cognitive lives of others? I owe the reader an answer to this question, but I have yet to offer one. In part, I think an answer to this question will turn on how we settle another thorny set of issues: what we are *licensed* to believe of each other. I admit this appears much like avoiding answering one difficult question by

suggesting that the answer turns on how we answer a different difficult question. Further, there is yet another difficult question we must answer before we can even address this challenge.

Who is it that we wrong and what kind of relationship do we need to stand in with them? Consider, for example, our beliefs about people in our distant past. When I see images of Roman graffiti and come to believe that whoever drew that was childish, have I come to wrong that Roman? Consider, also, Ea-Nasir. Ea-Nasir is a merchant from ancient Babylonia. We know about Ea-Nasir because a customer, Nanni, sent a letter—i.e., a clay tablet—complaining of the substandard copper that Ea-Nasir had sold him.[16] Must I take the participant stance towards Ea-Nasir (and Nanni)? The answer is clear in the case of deep personal relationships to each other, but when the relationship is thinner, indirect, or non-existent, to what extent can others have moral claims on our cognitive lives?

Further this dissertation has focussed only on the question of outright beliefs and how outright beliefs can wrong. But beliefs are not the only attitudes we hold towards each other. Further, neither are they the most common attitudes we hold towards each other. We not only believe things of others, we also hope, imagine, assume, expect, fantasize things about others. There is a whole range of belief-like attitudes we hold, and as a result, perhaps an even wider range of ways in which we might wrong each other. In some sense, I made the task for myself in this dissertation harder than it needed to be by focussing on belief. Instead, there could be an easier route to the thesis that beliefs can wrong by attending to the many belief-like attitudes we hold and how they wrong. For example, parents' expectations of their children and how their hopes for their children can wrong them. And if belief-like attitudes can wrong, perhaps it's a short step from the thesis that belief-like attitudes can wrong to the thesis that beliefs can wrong.[17] Further, our mental life is even richer than just belief and belief-like attitudes: these attitudes can also come in differing degrees of strength, i.e., credences. Can credences wrong in the same way outright beliefs wrong?

Finally, it seems that we not only wrong when we believe, we can also wrong when we refrain from belief and we can wrong when we undermine belief. Consider, for example, the Holocaust denier, the tobacco lobbyist insisting that the evidence is insufficient to support a link between smoking and lung cancer, the politician who denies that the evidence is sufficient to support anthropogenic climate change, the parent who denies that the evidence is sufficient to support the recommended vaccination schedule, etc. These

[16] See: https://en.wikipedia.org/wiki/Complaint_tablet_to_Ea-nasir

[17] Thanks to Amy Kind for pressing me on this point.

examples, I think, point to a broader phenomenon than what I have addressed in this dissertation: issues of epistemic injustice (see Fricker 2007).

Partly why I have not attended to the issue of epistemic injustice in this dissertation is because much of the current research on the topic of epistemic injustice has focused on cases involving testimony and questions of credibility assignment (Dotson 2011; Davis 2016; Pohlhaus 2017). In contrast, the kind of *wronging* I have been drawing attention to does not fall neatly into any of the discussion of epistemic injustice. That is, this kind of wronging involved in refraining from belief or undermining belief concerns not how an individual is wronged in their capacity as a knower, but in their capacity as a believer. I think this points to a phenomenon I very briefly gestured towards at the end of Chapter 3, *epistemic manipulation*. The Holocaust denier and the climate change skeptic both illustrate the way in which doubt can be harmful when it is unreasonable. However, just as it has been difficult to articulate necessary and sufficient conditions for when beliefs wrong, it can also be difficult to identify when doubt is unreasonable. This difficulty, however, it central to what makes the wrong so problematic.

Epistemic manipulation, as the counterpart to emotional abuse and emotional manipulation, can be difficult to identify because of how closely it resembles epistemic norms we standardly accept as good epistemic norms. For example, emotional abusers make use of norms that are part of a healthy relationship and they use those norms to manipulate their partners. Part of why it is so difficult to recognize emotional manipulation as a form of abuse is how closely it resembles healthy behaviour. Consider the following norm of a healthy relationship: you should trust your partner. An emotionally abusive partner, however, will manipulate that trust to harm their significant other.

An example of epistemic manipulation that has garnered some attention recently is gaslighting (see Abramson 2014; McKinnon 2017). Gaslighting is a form of emotional and epistemic manipulation that abuses the following standardly accepted epistemic norm: when someone you respect disagrees with you, many argue that you should revise your confidence in your initial belief. Consider the following case from the 1944 film, Gaslight. You believe that the lights have been flickering, but someone you respect repeatedly insists that they don't flicker. You are never able to prove that the lights flicker while this person is around. In normal circumstances, many argue that you should lower your confidence in the initial belief (Christensen 2007; Elga 2007; Feldman 2007). My hope is that the story I have told thus far in the dissertation can be extended to show how this case it is the epistemic norms themselves that are being used to wrong you. You are being epistemically manipulated.

In short, although I believe I have shown what I set out to show—that beliefs can

wrong and that when it comes to what we should (epistemically) believe morality is not voiceless—much more needs to be said about how the moral considerations of social injustice and epistemic norms interrelate. You might think, as Foot (2001) thinks about her own work at the end of *Natural Goodness*, that "[t]he proper reply is that in a way nothing is settled, but everything is left as it was." I agree that nothing is settled, but I also hope that it's not the case that everything is left as it was.

BIBLIOGRAPHY

BIBLIOGRAPHY

Abramson, K. (2014). Turning up the lights on gaslighting. *Philosophical Perspectives*, 28(1):1–30.

Antony, L. M. (2016). Bias: Friend or foe. In Brownstein, M. and Saul, J., editors, *Implicit Bias and Philosophy*, volume 1. Oxford University Press.

Appiah, K. A. (1990). Racisms. In Goldberg, D., editor, *Anatomy of Racism*, pages 3–17. University of Minnesota Press.

Appiah, K. A. (1995). The uncompleted argument: Dubois and the illusion of race. In Mosley, A. G., editor, *African Philosophy: Selected Readings*, pages 199–215. Englewood Cliffs.

Arpaly, N. (2004). *Unprincipled Virtue: An Inquiry into Moral Agency*. Oxford University Press.

Arpaly, N. and Schroeder, T. (2014). *In Praise of Desire*. Oxford University Press.

Atwood, M. (1998). *The Handmaid's Tale*. Anchor Books.

Baker, D. (MS). Skepticism about ought simplicter. *In progress*.

Banaji, M. and Greenwald, A. (2013). *Blindspot: Hidden Biases of Good People*. Delacorte Press.

Basu, R. and Schroeder, M. (2018). Doxastic wrongings. In Kim, B. and McGrath, M., editors, *Pragmatic Encroachment in Epistemology*. Routledge.

Beeghly, E. (2015). What is a stereotype? what is stereotyping? *Hypatia*, 30(4):675–691.

Begby, E. (2013). The epistemology of prejudice. *Thought*, pages 90–99.

Begby, E. (2018). Doxastic morality: A moderately skeptical perspective. *Philosophical Topics*.

Bell, M. (2009). Anger, virtue and oppression. In Tessman, L., editor, *Feminist Ethics and Social and Political Philosophy: Theorizing the Non-Ideal*. Springer.

Bero, S. (MS). Relationships and reactive attitudes. *in progress*.

Block, N. (1996). How heritability misleads about race. *The Boston Review*, 20(6):30–35.

Blum, L. (2002). Racism: What it is and what it isn't. *Studies in Philosophy and Education*, 21(3):204–218.

Blum, L. (2004). What do accounts of "racism" do? In Levine, M. and Pataki, T., editors, *Racism in Mind: Philosophical Explanations of Racism and Its Implications*. Cornell University Press.

Bolinger, R. (2017). *Mistaken Defense and Normative Conventions*. PhD thesis.

Bommarito, N. (2017). Virtuous and vicious anger. *Journal of Ethics and Social Philosophy*, 11(3).

Brink, D. (1994). Moral conflict and its structure. *The Philosophical Review*, 103:215–247.

Buchak, L. (2014). Belief, credence, and norms. *Philosophical Studies*, 169(2):285–311.

Burge, T. (1993). Content preservation. *The Philosophical Review*, 102.

Case, S. (2016). Normative pluralism worthy of the name is false. *Journal of Ethics and Social Philosophy*, 11(1).

Chang, R. (1997). *Incommensurability, Incompatibility and Practical Reason*. Harvard University Press.

Chang, R. (2003). All things considered. *Philosophical Perspectives*, 18(1):1–22.

Christensen, D. (2007). Epistemology of disagreement: The good news. *Philosophical Review*, 116(2):187–217.

Clifford, W. (1877). The ethics of belief. *Contemporary Review*.

Clough, S. and Loges, W. (2008). Racist value judgments as objectively false beliefs: A philosophical and social-psychological analysis. *Journal of Social Philosophy*, 39(1):77–95.

Conee, E. and Feldman, R. (2004). *Evidentialism*. Oxford University Press.

Copp, D. (1997). The ring of gyges: Overridingness and the unity of reason. *Social Philosophy and Policy*, 14(1):86–101.

Corlett, J. A. (2005). Race, racism, and reparation. *Journal of Social Philosophy*, 36(4):568–585.

Darwall, S. (2006). *The Second-Person Standpoint: Morality, Respect, and Accountability*. Harvard University Press.

Davis, E. (2016). Typecasts, tokens, and spokepersons: A case for credibility excess as testimonial injustice. *Hypatia*, 31(3):485–501.

de Beauvoir, S. (1949). *The Second Sex*, volume 2010. Random House.

DeRose, K. (1992). Contextualism and knowledge attributions. *Philosophy and Phenomenological Research*, 52:912–29.

DiAngelo, R. (2011). White fragility. *International Journal of Critical Pedagogy*, 3(3):54–70.

Dorsey, D. (2013). Two dualisms of practical reason. In Shafer-Landau, R., editor, *Oxford Studies in Metaethics*, volume 8. Oxford University Press.

Dotson, K. (2011). Tracking epistemic violence, tracking practices of silencing. *Hypatia*, 26(2):236–257.

Dotson, K. (2018). Accumulating epistemic power: A problem with epistemology. *Philosophical Topics*.

Doyle, A. C. (1887). *A Study in Scarlet*.

Doyle, A. C. (1892). The boscombe valley mystery.

Du Bois, W. E. B. (1903). *The Souls of Black Folk*, volume 1994. Dover.

Elga, A. (2007). Reflection and disagreement. *Nous*, 41(3):478–502.

Fantl and McGrath (2009). *Knowledge in an Uncertain World*. Oxford University Press.

Fantl, J. and McGrath, M. (2002). Evidence, pragmatics, and justification. *The Philosophical Review*, 111(1):67–94.

Fantl, J. and McGrath, M. (2007). On pragmatic encroachment in epistemology. *Philosophy and Phenomenological Research*, 75(3):558–589.

Feldman, R. (2004). The ethics of belef. In Conee, E. and Feldman, R., editors, *Evidentialism: Essays in Epistemology*, pages 166–195. Oxford University Press.

Feldman, R. (2007). Reasonable religious disagreements. In Antony, L., editor, *Philosophers without Gods*. Oxford University Press.

Finlay, S. (2006). The reasons that matter. *Australasian Journal of Philosophy*, 84(2):1–20.

Finlay, S. (2007). Too much morality. In Bloomfield, P., editor, *Morality and Self-Interest*. Oxford University Press.

Finlay, S. (2009). Oughts and ends. *Philosophical Studies*, 143(3):315–340.

Finlay, S. (2014). *Confusion of Tongues: A Theory of Normative Language*. Oxford University Press.

Flowerree, A. (2016). Agency of belief and intention. *Synthese*.

Foot, P. (1972). Morality as a system of hypothetical imperatives. *The Philosophical Review*, 81(3):305–316.

Foot, P. (2001). *Natural Goodness*. Oxford University Press.

Fricker, M. (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford University Press.

Fritz, J. (2018). From pragmatic encroachment to moral encroachment. *Pacific Philosophical Quarterly*.

Ganson, D. (2008). Evidentialism and pragmatic constraints on outright belief. *Philosophical Studies*, 139(3):441–458.

Garcia, J. L. A. (1996). The heart of racism. *The Journal of Social Philosophy*, 27(1):5–45.

Garcia, J. L. A. (1997). Current conceptions of racism: A critical examination of some recent social philosophy. *Journal of Social Philosophy*, 28(2):5–42.

Garcia, J. L. A. (1999). Philosophical analysis and the moral concept of racism. *Philosophy and Social Criticism*, 25(5):1–32.

Gendler, T. S. (2011). On the epistemic costs of implicit bias. *Philosophical Studies*, 156:33–63.

Gordon, L. (1995). *Bad Faith and Antiblack Racism*. Humanities Press.

Gordon, L. (2000). *Existentia Africana: Understanding Africana Existential Thought*. Routledge.

Greenwald, A. and Banaji, M. (2015). Statistically small effects of the implicit association test can have societally large effects. *Journal of Personality and Social Psychology*, 108(4):553–561.

Grimm, S. R. (2016). Knowledge, practical interests, and rising tides. In Greco, J. and Henderson, D., editors, *Epistemic Evaluation: Purposeful Epistemology*. Oxford University Press.

Hacking, I. (1995). The looping effects of human kinds. In Dan Sperber, D. P. and Premack, A. J., editors, *Causal cognition: A multidisplinary debate*, pages 351–394. Oxford University Press.

Hacking, I. (1999). *The social construction of what?* Harvard University Press.

Haslanger, S. (2012). *Resisting Reality: Social Construction and Social Critique*. Oxford University Press.

Hawthorne, J. (2003). *Knowledge and Lotteries*. Oxford University Press.

Hawthorne, J. and Stanley, J. (2008). Knowledge and action. *Journal of Philosophy*, 105(10):571–590.

Herrnstein, R. and Murray, C. (1996). *Bell Curve: Intelligence and Class Structure in American Life*. Free Press.

Hieronymi, P. (2005). The wrong kind of reason. *Journal of Philosophy*, 102(9).

Hieronymi, P. (2006). Controlling attitudes. *Pacific Philosophical Quarterly*, 87:45–74.

Hieronymi, P. (2008). Responsibility for believing. *Synthese*, 161:357–373.

Horty, J. F. (2003). Reasoning with moral conflicts. *Nous*, 37(4):557–605.

Ikuenobe, P. (2011). Conceptualizing racism and its subtle forms. *Journal for the Theory of Social Behaviour*, 41(2).

James, W. (1896). The will to believe. *The New World*.

Joyce, R. (2001). *The Myth of Morality*. Cambridge University Press.

Kant, I. (1785). *Groundwork of the Metaphysics of Morals*, volume 2012 of *Cambridge Texts in the History of Philosophy*. Cambridge University Press.

Kelly, T. (2003). Epistemic rationality as instrumental rationality: A critique. *Philosophy and Phenomenological Research*, 66(3).

Langton, R. (1992). Duty and desolation. *Philosophy*, 67.

Lengbeyer, L. (2004). Racism and impure hearts. In Levine, M. and Pataki, T., editors, *Racism in Mind: Philosophical Explanations of Racism and Its Implications*. Cornell University Press.

Leslie, S.-J. (2015). The original sin of cognition: Fear, prejudice, and generalization. *Journal of Philosophy*.

Locke, J. (1975). *An Essay Concerning Human Understanding*. Oxford University Press.

Lorde, A. (1984). Eye to eye: Black women, hatred, and anger. In *Sister Outsider*, pages 145–75. Crossing Press.

Lynn, M. (2006). Race differences in restaurant tipping: A literature review and discussion of practical implications. *Journal of Foodservice Business Research*, 9(4).

Madva, A. and Brownstein, M. (2016). Stereotypes, prejudice, and the taxonomy of the implicit social mind. *Nous*.

Mandelbaum, E. (2016). Attitude, inference, association: On the propositional structure of implicit bias. *Nous*, 50(3):629–658.

Marusic, B. (2012). Belief and difficult action. *Philosopher's Imprint*, 12(18).

Marusic, B. (2015). *Evidence and Agency*. Oxford University Press.

Marusic, B. and White, S. (2018). How beliefs can wrong—a strawsonian epistemology. *Philosophical Topics*.

McGarty, C., Yzerbyt, V., and Spears, R. (2002). Social, cultural and cognitive factors in stereotype formation. In McGarty, C., Yzerbyt, V., and Spears, R., editors, *Stereotypes as Explanations*. Cambridge University Press.

McHugh, C. (2012). Epistemic deontology and voluntariness. *Erkenntnis*, 77:65–94.

McKinnon, R. (2017). Allies behaving badly: Gaslighting as epistemic injustice. In *The Routledge Handbook of Epistemic Injustice*. Routledge.

McLeod, O. (2001). Just plain 'ought'. *The Journal of Ethics*, 5:269–91.

Melton, D. H. (2009). The vulnerable self: Enabling the recognition of racial inequality. In Tessman, L., editor, *Feminist Ethics and Social and Political Philosophy: Theorizing the Non-Ideal*, pages 149–164. Springer.

Mills, C. (2007). White ignorance. In Sullivam, S. and Tuana, N., editors, *Race and Epistemologies of Ignorance*. SUNY Press.

Mills, C. W. (1997). *The Racial Contract*. Cornell University Press.

Mitova, V. (2008). Why w. k. clifford was a closet pragmatist. *Philosophical Papers*, 37(3):471–489.

Moss, S. (2018). *Probabilistic Knowledge*. Oxford University Press.

Moss, S. (MS). Moral encroachment. *in progres*.

Munton, J. (MS). The epistemic flaw with accurate statistical generalizations. *in progress*.

Nagel, T. (1970). Death. *Nous*.

Nelson, M. T. (2010). We have no positive epistemic duties. *Mind*, 119(473).

Nussbaum, M. C. (2016). *Anger and Forgiveness*. Oxford University Press.

Oswald, F., Mitchell, G., Blanton, H., Jaccard, J., and Tetlock, P. (2013). Predicting ethnic and racial discrimination: A meta-analysis of iat criterion studies. *Journal of Personality and Social Psychology*, 105(2):171–192.

Oswald, F., Mitchell, G., Blanton, H., and Tetlock, P. (2015). Using the iat to predict ethnic and racial discrimination: Small effect sizes of unknown societal significance. *Journal of Personality and Social Psychology*, 108(4):562–571.

Pace, M. (2011). The epistemic value of moral considerations: Justification, moral encroachment, and james' 'will to believe'. *Nous*, 45(2):239–268.

Parfit, D. (2011). *On What Matters*. Oxford University Press.

Parker, P. (1978). *Movement in Black*, volume http://condor.depaul.edu/mwilson/multicult/patparker.htm, chapter For the white person who wants to know how to be my friend.

Paul, S. and Morton, J. (2018). Believing in others. *Philosophical Topics*, 46(1).

Pohlhaus, G. (2017). Varieties of epistemic injustice. In *The Routledge Handbook of Epistemic Injustice*. Routledge.

Preston-Roedder, R. (2013). Faith in humanity. *Philosophy and Phenomenological Research*, 87(3):664–687.

Preston-Roedder, R. (2017). Civic trust. *Philosophers' Imprint*, 17(4).

Preston-Roedder, R. (2018). Three varieties of faith. *Philosophical Topics*.

Rabinowicz, W. and Ronnow-Rasmussen, T. (2004). The strike of the demon: On fitting pro-attitudes and value. *Ethics*, 114.

Ray, K. (2017). Intersectionality and the dangers of white empathy when treating black patients. http://www.bioethics.net/2017/03/intersectionality-and-the-dangers-of-white-empathy-when-treating-black-patients/.

Ross, J. and Schroeder, M. (2012). Belief, credence, and pragmatic encroachment. *Philosophy and Phenomenological Research*.

Sartre, J.-P. (1943). *Being and Nothingness: An Essay on Phenomenological Ontology*, volume 2013. Routledge.

Scanlon, T. (1998). *What We Owe to Each Other*. Harvard University Press.

Schauer, F. F. (2006). *Profiles, probabilities, and stereotypes*. Harvard University Press.

Schroeder, M. (2012a). Stakes, withholding, and pragmatic encroachment on knowledge. *Philosophical Studies*, 160(2).

Schroeder, M. (2012b). The ubiquity of state-given reasons. *Ethics*, 122(3):457–488.

Schroeder, M. (2018a). Persons as things. In *Oxford Studies in Normative Ethics*. Oxford University Press.

Schroeder, M. (2018b). When beliefs wrong. *Philosophical Topics*, 46(1).

Shah, N. (2002). Clearing space for doxastic voluntarism. *The Monist*, 85(3):436–445.

Shelby, T. (2003). Review of "i'm not racist, but ...": The moral quandary of race by lawrence blum. *The Philosophical Review*, 112(1):124–6.

Shelby, T. (2016). *Dark Ghettos*. Harvard University Press.

Srinivasan, A. (2016). Would politics be better off without anger. https://www.thenation.com/article/a-righteous-fury/.

Srinivasan, A. (2017). The aptness of anger. *Journal of Political Philosophy*.

Srinivasan, A. (2018). Does anyone have the right to sex? *London Review of Books*, 40(6):https://www.lrb.co.uk/v40/n06/amia-srinivasan/does-anyone-have-the-right-to-sex.

Stanley, J. (2005). *Knowledge and Practical Interests*. Oxford University Press.

Steup, M. (2000). Doxastic voluntarism and epistemic deontology. *Acta Analytica*, 15(1):25–56.

Strawson, P. (1962). Freedom and resentment. In *Freedom and Resentment and Other Essays*, volume 2008. Routledge.

Stroud, S. (1998). Moral overridingness and moral theory. *Pacific Philosophical Quarterly*, 79(2):170–89.

Stroud, S. (2006). Epistemic partiality in friendship. *Ethics*, 116(3).

Sussman, D. (2005). What's wrong with torture? *Philosophy Public Affairs*, 33(1).

Thomson, J. J. (2001). *Goodness and Advice*. Princeton University Press.

Thomson, J. J. (2008). *Normativity*. Open Court.

Tiffany, E. (2007). Deflationary normative pluralism. *Canadian Journal of Philosophy*, 37(5):231–62.

Valian, V. (2005). Beyond gender schemas: Improving the advancement of women in academia. *Hypatia*, 20(3).

Velleman, J. D. (2005). A brief introduction to kantian ethics. In *Self to Self: Selected Essays*, pages 16–44. Cambridge University Press.

Wallace, R. J. (1994). *Freedom and the Moral Sentiments*.

Wallace, R. J. (2014). Emotions and relationships: On a theme from strawson. In *Oxford Studies in Agency and Responsibility*, volume 2, pages 119–42.

Weatherson, B. (2005). Can we do without pragmatic encroachment? *Philosophical Perspectives*, 19(1):417–443.

Weatherson, B. (2008). Deontology and descartes' demon. *Journal of Philosophy*, 105:540–569.

Weatherson, B. (2012). Knowledge, bets, and interests. In Brown, J. and Gerken, M., editors, *Knowledge Ascriptions*, pages 75–103. Oxford University Press.

Wedgwood, R. (2007). *The Nature of Normativity*. Oxford University Press.

Wedgwood, R. (MS). The unity of normativity. *in progress*.

White, R. (2005). Epistemic permissiveness. *Philosophical Perspectives*, 24:573–615.

Williams, B. (1965). Ethical consistency. *Proceedings of the Aristotelian Society*, 39.

Williams, P. J. (1992). *The Alchemy of Race and Rights*. Harvard University Press.

Williamson, T. (2000). *Knowledge and its Limits*. Oxford University Press.

Worsnip, A. (2015). Two kinds of stakes. *Pacific Philosophical Quarterly*, 96:307–324.