

The Wrongs of Racist Beliefs¹

Rima Basu

Forthcoming in *Philosophical Studies*.

We care not only about how people treat us, but also what they believe of us. If I believe that you're a bad tipper given your race, I've wronged you. But, 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 paper, I present 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, 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 paper I argue that although these accounts will capture many of the wrongs associated with racist beliefs, they will 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. To address this challenge, 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.

1. The Motivating Challenge

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

¹ I owe a large debt of gratitude to (in alphabetical order) the following people for lots of enlightening comments and discussions: Mike Ashfield, Renee Bolinger, Kenny Easwaran, Maegan Fairchild, Stephen Finlay, Liz Harman, Pamela Hieronymi, Gabbrielle Johnson, Tanya Kostochka, Zi Lin, Jessie Munton, Mark Schroeder, Ralph Wedgwood, various other graduate students and faculty at USC, several anonymous reviewers, and audience members at Athena in Action and the Penn-Rutgers-Princeton Social Epistemology Workshop where this paper was presented under the title “Motivating Ethical Demands on Belief”.

world.² 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 paper 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. To see this, consider the following case that is the motivating puzzle of this paper.

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 whenever I believe of a black diner in my section that they will tip worse than the 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 Lynn (2006), there is evidence of racial disparities in tipping practices that support Spencer's belief that the black diner in his section won't tip well. 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 (see also Gendler 2011).

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 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

² 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>

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 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 that Spencer is not committing a wrong. Sometimes there are unpleasant things that we ought believe. For example, it might be unpleasant for me to believe that one day soon my favourite television show will come to an end (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.

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 if it 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 seems to commit 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 wrong by believing what he does.

You may, however, still push back and ask why we should think that Spencer's belief is racist. What is so racist about believing that black diners won't tip well? To answer this, I must take a step back. The task of the paper is to take a paradigmatically racist belief, e.g., black diners do not tip well, and see what resources a supposedly rational racist can muster in defense of the epistemic grounds of such a belief in order to demonstrate that even the best argument for how the belief might be rationally held does not succeed.⁴

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 4.

⁴ I have made the conscious choice to use this belief, one regarding tipping practices, over others because (i) it is commonly accepted as being racist, and (ii) it doesn't require invoking the more offensive claims about black

There is still further question of explaining which beliefs wrong others. That is, the job fully of accounting the wrongs of racist beliefs is incomplete until we have a clear statement of the way in which the beliefs that are racist are racist. This is a large question and one that cannot be adequately addressed in this paper. The goal of this paper, despite the grandiose title, is more modest. The goal of this paper is to just make space for the idea that beliefs can wrong. The next task is to then explain which beliefs wrong. Just as we recognize that we can wrong others through the acts we commit but that not all acts wrong, we can recognize that there should be a way to delineate beliefs with racial content that are *racist and wrong* and beliefs *with racial content that are not racist and do not wrong*. For the time being, however, I must just ask the reader to follow me in granting that the belief that Spencer holds at the very least seems racist and seems to wrong. With that granted, we can explore where that wrong should be located.

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*. 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 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. This will be the guiding assumption as we proceed.

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

Our task is to explain the wrong Spencer commits. 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

people that are often invoked that are problematic in their own right. In short, I cannot answer why we should think that Spencer's belief is racist if you do not already accept (or are willing to grant) that it is a racist belief. To explain why it is racist would lead to a further question that cannot be adequately addressed in the space of this paper.

⁵ You might worry, as an anonymous referee has worried, that that this constraint seems to conflate a theoretical task – explaining what makes Spencer's belief racist and what makes it wrong – with the pragmatic task of engaging fruitfully with Spencer. I do not have a direct argument that Spencer's self-conception is correct, instead what I hope to provide is a suggestive argument that it would be disappointing if there could be no wrong if Spencer's self-conception is correct. If that is how things turn out, then we although we might have all the theoretical resources we need to diagnose real-life Spencers whose self-conceptions are not completely accurate, we cannot engage with them. They, after all, think they are in a better position than us to know whether they have been dispassionately rational or biased in forming their belief. It is thus helpful to have this pragmatic constraint guiding the theoretical task of explaining what makes Spencer's belief racist. It may, nonetheless, turn out that the right answer to the theoretical question—what makes Spencer's belief wrong—involves attributing to Spencer an epistemic mistake or a moral failing, and as a result, we shouldn't grant Spencer his self-conception of himself. Nonetheless, the pragmatic constraint that we should, as much as possible, do our best to grant Spencer his conception of himself should constrain how we approach the theoretical question.

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 according to the downstream theorist 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 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 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 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 Blum’s (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 *racial* 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.

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.

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 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 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

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.

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 (nor cannot) wrong Sanjeev because Sanjeev will never learn of the hermit's belief. However, neither harm nor knowledge are a prerequisite for being 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.

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 is right that people wrong each other in virtue of what they believe about each other insofar as 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.

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. *hurt*. For more on these cases, see Basu (2018, chapter 2).

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 state 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




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. So, 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

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

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. Leslie (2015) for example, argues that stereotypes can problematically lead to the possession of generic beliefs if that group is perceived to share an underlying *essence*. For Leslie 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, and in turn beliefs with generic content, when they are negative. In sum, the stereotype itself may be unproblematic. There could, perhaps, be creatures who could unproblematically hold stereotypes concerning groups and do nothing wrong in doing so.¹¹ But given our own cognitive limitations and predispositions, stereotypes wrong. Further, it is not the belief in the stereotype that 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

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

¹¹ For example, I can unproblematically hold as my stereotype of a dog the image of Labrador retriever, but it doesn’t follow from my possession of that stereotype that I also accept the generic claim that “Dogs are Labrador retrievers”. Thanks again to an anonymous referee for this point.

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.

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 that 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 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 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 needed) 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

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 Gabrielle Johnson for many conversations about implicit bias.

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 Muntun (MS). Begby argues that the reason why prejudices are epistemically insidious concerns how they become internalized as background beliefs that are recalcitrant to empirical counterevidence. Muntun 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

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, 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 Appiah (1990, 1995). Appiah 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 racist belief, so one cannot have racial prejudice without accepting racist beliefs.¹⁵ The epistemic wrong of racism, then, is a false belief about essences. The moral 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 treating some people morally.

Further, akin to the upstream accounts discussed in Section 3, part of the wrong of racism comes from the racist disposition that imbues in the believer a tendency to accept, morally

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 paper.

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.

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, 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 Garcia's (1996, 1997, 1999) 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's (2014) conception of the wrong of prejudice more generally: 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 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-semitic who has the false belief *Jews are members of a worldwide conspiracy set on world domination*. Arpaly argues that the anti-semitic 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 people of Jewish descent 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 people of Jewish descent 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

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-semitic. 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.

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 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 overreliance on the hot condition can obscure other cases of beliefs that wrong. For example, consider Valian (2005)'s account of the gender disparity in philosophy, which she calls a "cold" social-cognitive account given that it concerns the beliefs we come to hold on the basis of evidence that seems well-founded given our environment, e.g., assuming that the man at a table is more likely to be the CEO than a woman given that a disproportionate number of CEOs are men. As she notes, her account is "intended to explain what goes wrong in environments where nothing seems to be wrong, where people genuinely and sincerely espouse egalitarian beliefs and are well-intentioned, where few men or women overtly harass women" (Valian 2005, pp. 1). Furthermore, 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 (see Bonilla-Silva 2009). The hot condition, then, is not necessary for identifying the moral wrong of racism. 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.

5. A Final Hurdle: A Worry About Control

My aim in this paper 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 paper 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 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.

There is a final hurdle to clear to establish the plausibility 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 will 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, 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. Steup (2000) similarly argues that our beliefs are responsive to epistemic reasons and in this way we exercise control over our beliefs. Further, Weatherson (2008) has persuasively argued that this argument relies on too narrow a conception of voluntary, and 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, 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. 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 (for more on this point, see Basu and Schroeder Forthcoming).

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.

Works Cited

- 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.
- Banaji, M. and Greenwald, A. (2013). *Blindspot: Hidden Biases of Good People*. Delacorte Press.
- Basu, R. (2018). *Beliefs That Wrong*. PhD Dissertation. University of Southern California.
- Basu, R. and Schroeder, M. (Forthcoming). Doxastic Wronging. In *Pragmatic Encroachment in Epistemology*, eds. Brian Kim and Matthew McGrath. 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.
- 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.
- Bonilla-Silva, E. 2009. *Racism without Racists: Color-Blind Racism and the Persistence of Racial Inequality in America*, 3rd edition. Rowman & Littlefield.

- 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.
- Corlett, J. A. (2005). Race, racism, and reparation. *Journal of Social Philosophy*, 36(4):568– 585.
- Flowerree, A. (2016). Agency of belief and intention. *Synthese*.
- 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. (2011). On the epistemic costs of implicit bias. *Philosophical Studies* 156(1):33-63.
- Gordon, L. (1995). *Bad Faith and Antiracist Racism*. Humanities Press.
- Gordon, L. (2000). *Existential 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.
- Hieronimi, P. (2006). Controlling attitudes. *Pacific Philosophical Quarterly*, 87:45–74.
- Hieronimi, P. (2008). Responsibility for believing. *Synthese*, 161:357–373.
- Ikuenobe, P. (2011). Conceptualizing racism and its subtle forms. *Journal for the Theory of Social Behaviour*, 41(2).
- 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*.
- 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.

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.

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

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

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.

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.

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

Valian, V. (2005). Beyond gender schemas: Improving the advancement of women in academia. *Hypatia*, 20(3).

Weatherson, B. (2008). Deontology and Descartes' Demon. *Journal of Philosophy*, 105:540–569.