Why Are Some Statistical Generalizations Epistemically Risky?

Maeve Marley

Thesis submitted to the faculty of the Virginia Polytechnic Institute and State University
in partial fulfillment of the requirements for the degree of

Master of Arts
In
Philosophy

Jordan L. MacKenzie
Wendy S. Parker
Mercedes M.Corredor

April 11th, 2023
Blacksburg, VA

Keywords: Moral Encroachment, Epistemology, Epistemic Injustice, Statistical Generalizations

Why Are Some Statistical Generalizations Epistemically Risky?

Maeve Marley

ABSTRACT

Moral encroachment theses (MET) operate like pragmatic encroachment theses. When the stakes of belief are high, so are the standards for evidence. This means that evidence which is sufficient in a low stakes-of-belief scenario may be insufficient when the stakes are raised. Simply, METs aim to appeal to the varying moral intuitions that one may have in cases with different moral stakes and build an epistemological difference out of that moral distinction. For example, one might think that in cases of racial profiling, because the moral stakes of belief are high, what would otherwise constitute good evidence for belief is insufficient.

However, most METs assume that the probabilistic evidence on which one relies to form their belief is good evidence. Instead of examining the reliability of statistical generalizations, like those used in cases of racial profiling, the moral encroacher focuses on the moral facts of the circumstance of belief formation to explain why the subsequent belief is wrong *epistemically*. I will focus on Sarah Moss's account because she focuses on cases in which one forms an opinion on the basis of probabilistic evidence. I use Moss's version of the MET as a target to illustrate the challenges METs face in general. Broadly, Moss holds that a judgment's *moral risk* bears on its epistemic status.

In *Section 1,* I briefly outline Sarah Moss's MET and explain why it fails to identify which cases produce epistemically problematic judgments and fails to explain why those judgments are epistemically problematic. In *Section 2*, I offer an alternative account, which explains why statistical generalizations about marginalized social groups are likely unreliable as evidence. Thus, use of this kind of evidence leads to epistemically problematic beliefs. I conclude by introducing *epistemic risk* as an explanation for why the inference made in *Shopper* is epistemically problematic while the inference made in *Fraternity Member* is not.

Why Are Some Statistical Generalizations Epistemically Risky?

Maeve Marley

GENERAL AUDIENCE ABSTRACT

Imagine a shopkeeper who has just realized something was stolen from his shop. There are two possible suspects: a young white man and a young Black man. He did not see the shoplifting occur, and the only evidence he has is the statistical evidence *that young Black men are 70% more likely to shoplift than young white men*. By all accounts, he is not racially biased, this is simply a statistical fact that he is aware of. Based on this evidence, he forms the judgment that the young Black man is the likely culprit. Let's call this case *Shopper*.[1]

Now imagine a student on a college campus whose friend has been assaulted. There are two possible suspects: a young man who is not a fraternity member and a young man who is in a fraternity. The only evidence that the student has is the statistical evidence *that men involved in fraternities are 70% more likely to have committed sexual violence than average.* By all accounts she is not anti-fraternity, she is simply aware of this statistical evidence. Based on this evidence, she forms the judgment that the fraternity member is the likely assailant. Let's call this case *Fraternity Member*.

I think there's a difference between these two cases. Specifically, I think it's okay to make the inference in the latter case, but not in the former.[2] Even if you don't quite share my intuition, you might still think that however 'icky' it feels to draw the above sort of inference in *Fraternity Member*, it feels ickier still to draw it in *Shopper*. Either way, I don't think these intuitions are merely responsive to the moral facts of the cases: I think there's something different about the evidence relied upon in these cases. Specifically, we have reason to thinks that the processes with which we produce the evidence relied upon in *Shopper* are biased.

---

[1] This example is adapted from Moss (2018) which Moss based on Beaty (2014).

[2] As outlined by Romeijn (2022): For some statisticians, it would be a mistake to make this kind of inference to probability or likelihood in the first place, since a given person either did or did not perform the action, but I am bracketing those concerns here.

Table of Contents

Introduction

Imagine a shopkeeper who has just realized something was stolen from his shop. There are two possible suspects: a young white man and a young Black man. He did not see the shoplifting occur, and the only evidence he has is the statistical evidence *that young Black men are 70% more likely to shoplift than young white men*. By all accounts, he is not racially biased, this is simply a statistical fact that he is aware of. Based on this evidence, he forms the judgment that the young Black man is the likely culprit. Let's call this case *Shopper*.[3]

Now imagine a student on a college campus whose friend has been assaulted. There are two possible suspects: a young man who is not a fraternity member and a young man who is in a fraternity. The only evidence that the student has is the statistical evidence *that men involved in fraternities are 70% more likely to have committed sexual violence than average.* By all accounts she is not anti-fraternity, she is simply aware of this statistical evidence. Based on this evidence, she forms the judgment that the fraternity member is the likely assailant. Let's call this case *Fraternity Member*.

I think there's a difference between these two cases. Specifically, I think it's okay to make the inference in the latter case, but not in the former.[4] Even if you don't quite share my intuition, you might still think that however 'icky' it feels to draw the above sort of inference in *Fraternity Member*, it feels ickier still to draw it in *Shopper*. Either way, I don't think these intuitions are merely responsive to the moral facts of the cases: I think there's something epistemically different about these two cases.

---

[3] This example is adapted from Moss (2018) which Moss based on Beaty (2014).

[4] As outlined by Romeijn (2022): For some statisticians, it would be a mistake to make this kind of inference to probability or likelihood in the first place, since a given person either did or did not perform the action, but I am bracketing those concerns here.

So, how do we justify these intuitions? When is it epistemically problematic to form beliefs about individual members of social groups based on statistical generalizations?

One answer to this question appeals to the varying moral intuitions that one may have in the two cases and tries to build an epistemological difference out of that moral distinction. This is the aim of the moral encroachment thesis (MET).[5] There are various ways of fleshing this thesis out, but in general METs seek to explain how moral facts bear on the epistemic status of beliefs.

Most METs assume that in cases like *Shopper* and *Fraternity Member*, the probabilistic evidence on which one relies to form their belief is good evidence. Instead of examining the reliability of statistical generalizations, like those used in cases of racial profiling, the moral encroacher focuses on the moral facts of the circumstance of belief formation to explain why the subsequent belief is wrong *epistemically*. I will focus on Sarah Moss's account because she focuses on cases in which one forms an opinion on the basis of probabilistic evidence. I use Moss's version of the MET as a target to illustrate the challenges METs face in general. Broadly, Moss holds that a judgment's *moral risk* bears on its epistemic status.

I find Moss's moral encroachment thesis unsatisfying in two respects. First, it seems we can have knowledge in morally risky cases. Second, it fails to give an adequate explanation of why even those cases that it successfully identifies as morally risky are thus epistemically problematic. Instead of looking at moral risk to explain why a belief

[5] In Bolinger (2020), they outline a variety of moral encroachment thesis and hold that in general, these theses hold that moral reasons can encroach on one's epistemic reasons for belief. In other words, "high moral stakes affect what it takes for a belief to be epistemically justified, or to count as knowledge" (p. 1). In general, moral encroachment theses hold that certain moral facts bear on epistemic facts and that one ought to be sensitive to these moral considerations when making epistemic judgements.

may be epistemically problematic, in what follows, I propose that we explore the reliability of statistical generalizations about different types of social groups as evidence. In the specific case of racialized groups, statistical generalizations often reflect epistemic injustices and may plausibly be taken to be unreliable evidence. I think this is also the case for marginalized social groups more generally. Relying on this evidence is *epistemically risky* as it promotes the use of unjust epistemic practices. Ultimately, it is the epistemic risk involved in *Shopper* that makes the shopkeepers belief epistemically problematic.

In *Section 1,* I briefly outline Sarah Moss's MET and explain why it fails to identify which cases produce epistemically problematic judgments and fails to explain why those judgments are epistemically problematic. In *Section 2*, I offer an alternative account, which explains why statistical generalizations about marginalized social groups are likely unreliable as evidence. Thus, use of this kind of evidence leads to epistemically problematic beliefs. I conclude by introducing *epistemic risk* as an explanation for why the inference made in *Shopper* is epistemically problematic while the inference made in *Fraternity Member* is not.

Moral Risk & Epistemically Problematic Beliefs

Moral encroachment theses operate like pragmatic encroachment theses. When the stakes of belief are high, so are the standards for evidence. This means that evidence which is sufficient in a low stakes-of-belief scenario may be insufficient when the stakes are raised. Consider the following case as an example of pragmatic encroachment. Imagine that you have just left home for the day. You can picture yourself locking the

door as you left, so you believe that the door is in fact locked. You live in a safe neighborhood, so the stakes of your belief are relatively low. Now, imagine instead that you live in a dangerous neighborhood. Again, you can picture yourself locking the door, but the stakes are high (if your door is unlocked valuable possessions may be stolen). Consequently, you believe that the door is not locked. Though your evidence in both cases is the same, pragmatic encroachment theses hold that it is reasonable for our standards of evidence to shift relative to the stakes of a given situation.[6] Simply, pragmatic considerations bear on the epistemic status of our judgements.

Similarly, moral encroachment theses hold that the moral features of a belief may impact its epistemic status. Different METs identify different relevant moral features. Sarah Moss' particular version of moral encroachment (M-MET) pinpoints *moral riskiness* as the moral feature of a judgment[7] which prevents that–otherwise justified—judgment from constituting knowledge. To illustrate Moss's thesis, imagine the following case:

> *Racial Profiling***:** A police officer is driving through a predominantly white neighborhood when she sees a young Black man. She knows that a crime has recently been committed in the area. She is aware of the statistical fact that young Black men are 70% more likely than young white men to be guilty of some crime. In response to this evidence, she judges that this young man is more likely to be guilty of something than other young people in the area.

According to Moss, because the officer's opinion is morally risky (i.e. the young man would be morally harmed if the officer were to falsely believe him to be guilty) the

[6] This example is adapted from Singh (unpublished).
[7] Moss (2018) uses the term *opinion.*

officer ought to consider other relevant possibilities beyond the statistical evidence she has.[8] When the officer does consider these possibilities, her larger body of evidence contains *probabilistic content*[9] which she cannot rule out, so her body of evidence becomes inconsistent. Thus, she is unable to form the judgment that the young Black man is more likely to be guilty. If the officer does not make this consideration, her judgment is epistemically problematic. Either way, the officer fails to *know* that the young man is likely to be guilty.[10]

Crucially, Moss's account holds that in all cases where the judgment formed by an epistemic agent is *morally* risky, it fails to constitute knowledge. Though Moss seems right that one ought to consider the other relevant possibilities to their known statistical evidence, it is unclear: 1) that in every morally risky case we fail to have knowledge and 2) how moral risk bears on a judgment's epistemic status. While her account vindicates opponents of racial profiling, it produces strange conclusions when applied to morally risky cases more generally. Consider the following case:

> *Children's Toys*: Susan works as a product designer for children's toys. Due to a process necessary for its construction, the particular toy she's working on contains a trace amount of a toxic chemical. However, after rigorous testing and

[8] Moss (2018) refers to these as 'probabilistic contents' which cannot be ruled out. She holds that in morally risky cases, a good moral agent ought to abide by the moral rule of consideration (ie. the officer ought to consider relevant alternatives like this particular young man does not behave like his general social group) and that this consideration does not establish the dependence of epistemic features on moral features (Moss, 191). However, regardless of whether this consideration of other possibilities is made, those possibilities have an epistemic impact which the M-MET reflects (Moss,192).

[9] This is a term borrowed from Moss (2018). She (and I) use it to refer to bits of probabilistic evidence which challenge the probabilistic evidence upon which one is relying to form the judgment in question. This can also be understood as the 'relevant possibilities' which contradict the probabilistic evidence on which a judgment relies.

[10] In this particular sort of case, Moss (2018) holds that while the officer may have a justified high credence that the young Black man is likely to be guilty of some crime, this fails to constitute knowledge (Moss, p. 9-10).

experimentation, it is highly unlikely that any toy contains a concentration of the toxic chemical high enough to cause children harm.

Though Susan is sure that the amount of the toxic chemical is so low that it is safe for children, her judgment is morally risky. In other words, if Susan were to be wrong, children could be seriously harmed. According to Moss's MET, this potential moral harm is just the kind of relevant moral feature that bears on the epistemic status of Susan's judgment. When applied to Susan's case, the M-MET holds that Susan cannot know that the amount of toxic chemical left in the toy is safe. To the particularly cautious consumer, or the concerned parent, this conclusion may seem reasonable. However, imagine that Susan has tested her toy again and again in every possible scenario of manufacturing and use. After each of these tests, the results are the same. Susan is certain that the amount of the toxic chemical in the toy is incredibly low, in fact its presence in this toy is lower than most other toys currently on the market. A mother herself, Susan is so confident in the toy's safety she has given it to her own children.

According to Moss, like in *Racial Profiling*, Susan ought to consider the other relevant possibilities to the toy being safe. Perhaps her tests were faulty or incomplete. When she considers these possibilities, Moss thinks that her body of evidence becomes inconsistent. By this Moss means that Susan has both probabilistic evidence of the toy's safety, and *probabilistic content*[11] which she cannot rule out. Thus, while she may have a justified high credence that the toy is safe, she lacks knowledge.

I think *Children's Toys* illustrates two ways in which Moss's MET goes wrong. First, it shows that the M-MET produces strange consequences when widely applied to

[11] In Susan's case, this is something like 'chemical testing is vulnerable to error, so this test may be faulty' or generally the relevant possibility that the toy may in fact be dangerous.

morally risky cases. Because Susan's case is morally risky, even after hundreds of rounds of testing, she cannot know the toy is safe. This claim seems too strong. Surely at some point Susan's testing is exhaustive and she is able to rule out those relevant possibilities which may contradict her evidence and know the toy is safe, despite the moral risk involved.

With this in mind, it is worth considering how the probabilistic contents which Susan cannot rule out ought to function as evidence. Perhaps Moss thinks these possibilities are a kind of higher order evidence against the reliability of her first order probabilistic evidence of the toy's safety. If this is the case, it seems as though this higher order evidence ought to be considered in all cases in which one relies on probabilistic evidence. However, Moss holds that it is due to the moral risk involved that Susan ought to consider the other relevant possibilities. So, the question remains: why is this consideration only necessary in morally risky cases?

This leads to the second way in which Moss's MET goes wrong. Specifically, it fails to explain how a judgment's moral risk bears on its epistemic status. Imagine that instead of testing children's toys, Susan is testing decking (*Decking*). Like in *Children's Toys*, the decking is made according to the same process and contains some amount of a toxic chemical. Susan tests and re-tests the decking and concludes that the remaining amount of toxic chemical in the decking is very low. According to Moss, because her judgment is not morally risky (i.e. no one is at risk of moral harm) she need not consider the relevant possibilities that her probabilistic evidence cannot fully rule out. So, in this case, Susan's body of evidence is not inconsistent. Thus, her justified high credence that

the decking contains a particular level of toxic chemical may constitute knowledge.[12] Perhaps Moss thinks that the consideration of such evidence makes for responsible epistemic practice. But, if this is the case, why is Susan permitted to be less epistemically responsible when testing decking? Assuming the probabilistic evidence is the same in both cases, it is not clear why moral risk matters.

Moss does not face this challenge alone. In general, METs rest on the contentious claim that a judgment's epistemic status may depend on its moral features. For the moral encroacher, some moral feature has to intervene on the epistemic judgment in order to reasonably distinguish cases like *Children's Toys* and *Decking*. This is in part because METs generally assume that the evidence on which one relies to form a given judgment is good evidence. Bolinger (2020) defines moral encroachment as: "Holding fixed the degree of probabilistic support the evidence gives for p, whether an epistemic attitude about p has some positive epistemic status can depend importantly on its moral features" (p.6). For example, the moral encroacher assumes that the evidence relied upon in both *Racial Profiling* and *Children's Toys* is reliable. As such, they locate the problematic nature of judgements made on the basis of this evidence downstream. By this I mean that they seek out moral features which may intervene on the epistemic status of certain judgements to explain why those judgements are problematic. This allows moral encroachers to conclude some judgements are epistemically problematic and others are not, even when the evidence in both cases is the same. I think this is a mistake. Though

[12] In this case, Susan has a justifiably high credence which can constitute knowledge. Because her judgment that the decking likely contains a particular level of toxic chemical is not morally risky, she need not consider other relevant possibilities and thus her body of evidence is consistent. According to Moss, this is sufficient for her having knowledge in this case (Moss, p. 9-10).

both cases are morally risky, the evidence relied upon in *Racial Profiling* is importantly different than that in *Children's Toys*.

Instead of citing a judgment's moral risk as explanation for its problematic epistemic status, I want to suggest that probabilistic evidence which aims to describe the behavior of marginalized social groups is very likely bad evidence. Further examining this evidence can help illuminate where the officer in *Racial Profiling* goes wrong. Unlike Susan, the officer in *Racial Profiling*, relies upon evidence heavily influenced by the complex social-epistemic terrain in which knowers of different social locations are more or less able to dictate what sorts of things count as reliable evidence and why. In other words, the intuition that the officer's conduct is not only morally but also epistemically problematic[13] may be explained by a closer examination of her evidence.

Further, my account will hold that not all cases of profiling based on statistical generalizations produce epistemically problematic judgments. Recall the two cases presented in the introduction. While both cases involve the formation of a judgment about an individual based on statistical evidence about the social group to which they belong, the inference formed in *Shopper* seems epistemically problematic in a way that the inference formed in *Fraternity* does not. My account justifies this intuitive asymmetry.

Kinds of Generalizations: What Makes for Unreliable Evidence?

Rather than locate the *epistemic* problem with beliefs formed in cases like racial profiling in some relevant *moral feature*, I propose that statistical generalizations about marginalized social groups carry significant risk of being unreliable pieces of evidence.

[13] Moss holds that epistemically problematic beliefs cannot constitute knowledge.

This is because we have good reason to think that the epistemic processes which generate this evidence are biased. Particularly, the evidence relied upon reflects the conditions of injustice under which it is collected. Thus, these statistical generalizations risk inaccurately describing the general behaviors of the social groups they purport to describe.

Recall *Racial Profiling* from *Section 1*. Remember that the police officer has formed the judgment that a young Black man is more likely to be guilty of some crime than other young men in the area on the basis of the following statistical evidence: 'young Black men are 70% more likely than other young men to be guilty of some crime'. Though this statistical generalization seems to support the officer's belief, more investigation is required to determine whether or not this evidence accurately describes the general behavior of young Black men relative to other young men before it is deemed reliable. Specifically, the social and epistemic dimensions of the interactions which produce this sort of evidence ought to be critically examined in order for us to determine its reliability.

To better understand how these kinds of statistical generalizations are produced and reinforced, let's imagine what might happen when an officer and a young Black man interact. Perhaps the officer assumes that the young man is guilty of that crime and is not impartial when she looks for evidence. This may occur consciously or unconsciously, but either way the evidence she selects, and the inferences she makes based on that evidence, is likely informed by some degree of implicit, or explicit, racial bias.[14] For example,

[14] Spencer, Charbonneau & Glaser (2016) underscores the influence of implicit bias in policing.

imagine that the young man is holding an object that the officer infers to be a weapon.[15] When the young man tries to advocate for his innocence, the officer may note his irregular speech patterns and improper grammar as indicative of a lack of education and possible propensity toward criminal behavior.[16] Though neither of these observations necessarily offer confirming evidence of the young man's guilt, the officer may treat this information as such.

Imagine further that this kind of scenario is not an isolated incident. In fact, it isn't. Many officers across the nation approach young Black men in this way.[17] As such, they disproportionately arrest young Black men and these men are convicted of crimes.[18] Perhaps some are guilty, but perhaps many are innocent. Either way, the officer's testimony in favor of their guilt is powerful. Importantly, aggregations of biased epistemic practices produce statistical generalizations like 'young Black men are 70% more likely to be guilty of some crime than young white men'.

We can use resources from the epistemic injustice literature[19] to make sense of what is going on when the officer observes the teenager and how this impacts the teenager's ability to effectively provide evidence. In general, people are taken more or less seriously as knowers due to their social location. This holds serious epistemic consequences. Agents in dominant social locations are more likely to have direct access

[15] Fridell (2016) found that "the police officers were better able to distinguish weapons from non-threatening objects, when held by Black and Latino targets than by Asian and White targets" (pp. 306–307).

[16] Brunson & Weitzer (2009) found that police were likely to distrust the testimony of young Black men due to their speech, dress and appearance.

[17] Brooks, M., Ward, C., Euring, M. et al. (2016) explains the history of criminalization of black men in America by police.

[18] Hinton, Henderson & Reed (2018) compile evidence which demonstrates young Black men are more likely to be arrested, incarcerated and given longer sentences.

[19] See: Haraway (1988), Crenshaw (1989), Harding, S. (1991), Hill, C. (1991), Fricker (1999, 2007), Pohlhaus (2011), Medina (2011),Dotson (2014) & Toole (2019).

to *shared epistemic resources*. These resources consist of terms which describe common behaviors and experiences. However, the more unique an experience is to a particular social group, the less likely that experience has a term widely understood by audience members outside of that group. Things like language barriers, cultural differences and marginalized experiences of the social world can prevent one from expanding and refining those shared epistemic resources on which we all rely. A *dominant social epistemic structure* arises.

Because individuals who hold marginalized social identities lack access to and influence over shared epistemic resources, the language necessary to describe their nuanced social experiences may not exist or be accepted by the dominant social epistemic framework. This is referred to as *hermeneutical injustice*. For example, the Black teenager may use words or phrases that, while commonly understood within his community, are not well understood beyond that community. So, the officer may not understand the young man when he tries to explain what he is doing and why. This may lead the officer to 1) assume that these phrases indicate his guilt, or 2) take his testimony to be unreliable.

The teenager's inability to effectively communicate evidence of his innocence may also be understood in terms of *testimonial injustice*. Because the officer occupies a socially dominant position, her testimony is taken to be more reliably truth tracking simply because she is regarded as holding more social power. Further, the officer's elevated social position allows them (and more broadly their social group) to have direct access to and even significant influence over shared epistemic resources and the dominant social epistemic structure. Conversely, the young man experiences a *credibility*

deficit.[20] This means that his testimony is taken to be less reliably truth tracking than it in fact is simply because he holds less social power. Therefore, the dominant social epistemic framework in which both the young man and the officer interact as epistemic agents fails to appropriately incorporate the experience of the young man. This is the case widely for agents who hold marginalized identities. Despite the fact that their epistemic perspective is truth tracking, their experiences fail to meet dominant standards of evidence or belief justification. As such, their testimony is not given due credit and those attributed a credibility excess retain greater access to the dominant social epistemic structure.

Importantly, we can imagine this scenario playing out hundreds of times with different officers and young men of color. These young men are disproportionately approached, arrested and convicted by law enforcement. Perhaps, in many cases, they are guilty of some crime; however, the conditions under which evidence of their potential guilt is gathered is wrought with implicit and explicit bias.[21] The epistemic injustice literature helps illuminate how exactly this is the case, as explained above. Young black men are assigned credibility deficits and are excluded by the dominant social epistemic structure. Conversely, police are awarded a credibility excess and are able to influence the dominant structure such that their evidence is privileged. When these instances are aggregated, we can reasonably expect that a resulting statistic like: 'young Black men are 70% more likely to be guilty of some crime than young white men' will overestimate the criminal activity of young Black men relative to other young men; it is not clearly reliable.

[20] Fricker (2007) coined *testimonial injustice*.
[21] Hinton, E., Henderson, L. & Reed, C. (2018)

Here, one might object that these sorts of interactions, even if they are plentiful, need not influence the production of the statistical evidence at issue. Thanks to the social sciences, there are well formed methodologies by which one can produce data about the in-general behavior of various social groups. However, though ideal research methodologies may well exist, the reality of racism (among other isms) in our society is inescapable. All epistemic agents, even researchers, interact, collect data and form judgements in a complex social-epistemic sphere. The social-epistemic topography which allows some agents better access to dominant resources and more credibility cannot be idealized away.

Unlike proponents of the MET, I argue that because we have good reason to think that statistical generalizations about the general behavior of marginalized social groups are formed according to biased epistemic practices, they are very likely to be unreliable. Specifically, these generalizations reflect instances of epistemic injustice. As I have argued, there are significant social and epistemic factors at play during interpersonal interactions between epistemic agents, particularly when they hold different social-epistemic positions. Aggregations of these interactions produce the kinds of statistical evidence that the officer, and those in cases of profiling an individual based on their membership to a marginalized social group, largely rely upon. So, these generalizations more accurately describe the experience of the dominantly positioned knower (in this case, the officer). When one uses unreliable evidence to draw an inference about a particular group member, the resulting judgment is epistemically problematic.

While my examples have highlighted racial profiling, I think that unjust epistemic practices are commonly reflected in probabilistic evidence which aims to describe the

behavior of marginalized social groups. Members of these groups are generally assigned credibility deficits, experience hermeneutical injustice and broadly lack access to and influence over dominant social epistemic structures. Consider the initial (and persistent) disbelief that women experience when reporting instances of sexual assault.[22] Due to patriarchal power structures, women are assigned a credibility deficit while men are assigned a credibility excess. Moreover, because the larger social epistemic structure in which these testimonies are exchanged is skewed in favor of those with more social epistemic power, women – and more specifically survivors of sexual violence — are less able to influence the social epistemic structure by which the reliability of their testimony is judged. Therefore, like in racial profiling cases, we have good reason to expect that statistical generalizations which aim to represent the general behavior of survivors obscure the social reality they aim to represent.[23]

To see how my view differentiates *Shopper* and *Fraternity Member*, let's return to and examine these cases. Recall in *Shopper*, a shopkeeper has just realized something has been stolen from his shop. There are two possible suspects: a young white man and a young Black man. He did not see the shoplifting occur, the only evidence he has is that *Black men are 70% more likely to shoplift than white men*. By all accounts, he is not racially biased, this is simply a statistical fact that he is aware of. Based on this evidence, he forms the judgment that a young Black man is the likely culprit.

Like the officer, the evidence on which the shopkeeper relies on reflects epistemically unjust practices. Simply, non-Black shoppers are likely caught shoplifting

[22] Antaki, C., Richardson, E., Stokoe, E. and Willott, S. (2015)

[23] On my view, statistical generalizations which aim to describe the frequency of sexual assault are likely under-descriptive of the actual social phenomenon. So, we may have reason to think that we should always believe women who claim to be survivors. This is a plausible implication of my view that I will not explore further in this paper.

less as their testimony is found to be more credible and aspects of their appearance or behavior may not be as readily taken as evidence of guilt. In *Shopper*, like in *Racial Profiling*, the shop-keeper relies on bad evidence to form his judgment that the young Black man is likely guilty. Thus, his judgment is epistemically problematic.

However, this is not the case in *Fraternity Member*. Unlike in *Shopper*, the fraternity member is not being profiled due to a socially marginalized identity he holds. Greek organizations are historically white, and their membership is reserved only for those wealthy enough to attend higher education and pay additional membership dues. As such, the fraternity member is awarded a credibility excess in comparison with their accusers in virtue of his dominant social epistemic position.[24] Further, while data on over policing in Black and Brown neighborhoods is readily available, there is no parallel recorded phenomenon when it comes to the 'over policing' of fraternities. Therefore, there is no comparable evidence or reason to believe that higher rates of sexual violence committed by fraternity members are due to over-zealous university disciplinary offices.[25] Because the statistical evidence relied upon in *Fraternity Member* is not likely to be influenced by unjust epistemic practices in the same way, the judgment made in *Fraternity Member* is not epistemically problematic like that made in *Shopper*.[26] If anything, the statistical evidence in this case under-represents fraternity member's likelihood to commit sexual violence.

---

[24] Chang (2014)

[25] If we were to start seeing data that fraternities were being over-policed, we might have good reason to think that the statistical evidence relied upon in *Fraternity Member* is likely unreliable like that in *Shopper* and *Racial Profiling.* As such, my view would conclude that the judgment produced in *Fraternity Member* is epistemically problematic like that in *Shopper* and *Racial Profiling*.

[26] Unlike the shopper, the fraternity member is profiled due to a group membership he has chosen. This signals that he is willingly associated with the stereotypes related to fraternity life. Conversely, the shopper is profiled based on his race, and associated stereotypes, which he cannot select for or against. I think this signals another important sense in which *Shopper* and *Fraternity Member* are disanalogous.

In sum, if one accepts that a piece of statistical evidence is significantly influenced by the unjust power dynamics at work within epistemic structures, then we have good reason to think this evidence is unreliable. Specifically, statistical evidence regarding generalized group behaviors relies on social facts about those groups as understood and endorsed by dominant social epistemic structures. As such, the means of production of this kind of evidence fails to appreciate the testimony and other epistemic resources presented by the marginalized group members it aims to generalize and it fails to reliably describe the general behaviors it purports to. When this unreliable evidence is used to justify an individual's judgments, those judgments are epistemically problematic.

Epistemic Risk

Rather than locate a belief's epistemic wrongness in its moral features, I have offered an epistemic account of why beliefs formed in cases like *Shopper* and *Racial Profiling* are epistemically problematic while the belief formed in *Fraternity Member* is not problematic in the same regard.[27] Broadly, my account holds that beliefs formed on the basis of evidence which is likely unreliable are epistemically problematic. Statistical evidence which describes the behavior of marginalized social groups is significantly influenced by repeated instantiations of epistemic injustice against members of such groups whose experiences and testimony fails to be adequately appreciated by the dominant social epistemic structure. This explains why the beliefs formed in *Shopper* and *Racial Profiling* are epistemically problematic, but perhaps there is still more to be said about the riskiness of these cases.

[27] Again, some statisticians think that inferences of this kind are mistaken. See footnote 2.

Both *Shopper* and *Racial Profiling* are cases where the belief formed is morally risky. However, as established in *Section 1*, a judgment's moral riskiness does not clearly bear on that judgment's epistemic status. Yet, these cases seem risky in some other important respect. By this I mean that it is important to consider the epistemic injustice that the use of bad evidence promotes. When individual agents rely on bad evidence and fail to consider how certain generalized 'facts' over-represent the epistemic standpoints of socially powerful individuals they 1) form skewed judgments about individuals using evidence which does not reliably describe the behavior of the marginalized social group to which they belong and 2) maintain epistemic standards which allow for the use of this bad evidence as adequate justification.[28] I think cases like these are best understood as (at least)[29] *epistemically risky*.[30]

Recall *Shopper*. In this case, when the shopkeeper relies on the bad statistical evidence at his disposal, he effectively perpetuates epistemic injustice against the Black teenager and his social group more broadly. Structurally, the shopkeeper reinforces an epistemic framework which fails to adequately include the experiences and testimony of marginalized knowers. Similarly, epistemic injustice is committed against the young woman whose testimony of assault is not given the credibility it is due, but not against the fraternity member who is assigned an excess of credibility when he testifies to his innocence.

With this in mind, my account may be used to identify which cases of belief formation are epistemically risky. Cases in which an epistemic agent is presented with

---

[28] Medina (2011)

[29] At worst, agents forming judgements in these cases are epistemically irresponsible or even epistemically vicious.

[30] While these cases are also (mostly) morally risky, I think we should focus on the sense in which they are epistemically risky to explain why these beliefs are bad qua normal standards for belief justification.

statistical evidence which is influenced by epistemic injustice to form a belief about an individual member of a social group risk perpetuating epistemic injustice. Further, when such evidence is relied upon, the resulting belief is epistemically problematic.

Though I have critiqued the moral encroachment view, there may be some welcome symmetry in this regard. Plausibly, lots of morally risky cases are also epistemically risky. Further, my account succeeds in finding the kinds of cases the proponent of MET takes issue with as epistemically problematic while maintaining that some morally risky judgments are epistemically permissible. Finally, my account avoids making the contentious claim that a judgment's moral features bear on its epistemic status.

References

Anderson, E. (2012) 'Epistemic Justice as a Virtue of Social Institutions', *Social Epistemology*: 26, 163-173.

Antaki, C., Richardson, E., Stokoe, E. and Willott, S. (2015), 'Police interviews with vulnerable people alleging sexual assault: Probing inconsistency and questioning conduct', *J Sociolinguistics*, 19: 328-350.

Bolinger, R.J. (2020), 'Varieties of Moral Encroachment', P*hilosophical Perspectives*: 34, 5-26.

Brooks, M., Ward, C., Euring, M. et al. (2016) 'Is There a Problem Officer? Exploring the Lived Experience of Black Men and Their Relationship with Law Enforcement' *Journal of African American Studies*: 20, 346–362.

Brunson, R., Weitzer, R. (2009) 'Police Relations with Black and White Youths in Different Urban Neighborhoods', *Urban Affairs Review*: 44(6), 858-885.

Crenshaw, K. (1989) 'Demarginalizing the Intersection of Race and Sex: A Black Feminist Critique of Antidiscrimination Doctrine, Feminist Theory and Antiracist Politics' *University of Chicago Legal Forum*: 1989, 139-167.

Dotson, K. (2014) 'Conceptualizing Epistemic Oppression', *Social Epistemology*: 28, 115-138.

Fricker, M. (1999) 'Epistemic Oppression and Epistemic Privilege' Canadian Journal of Philosophy: 25, 191-210.

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

Fridell, L. (2017). The Science of Implicit Bias and Implications for Policing. In: Producing Bias-Free Policing. SpringerBriefs in Criminology. Springer, Cham.

Haraway, D. (1988) 'Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective', *Feminist Studies:* 14 (3), 575–99.

Harding, S. (1991) 'Whose Science? Whose Knowledge?: Thinking from Women's Lives', *Cornell University Press.*

Hill, C. (1991) 'Black Feminist Thought: Knowledge, Consciousness, and the Politics of Empowerment', *Routledge: New York.*

Hinton, E., Henderson, L. & Reed, C. (2018) 'An Unjust Burden: The Disparate Treatment of Black Americans in the Criminal Justice System', *Vera Institute of Justice*: 1-20

Romeijn, Jan-Willem, "Philosophy of Statistics", The Stanford Encyclopedia of Philosophy (Fall 2022 Edition), Edward N. Zalta & Uri Nodelman (eds.).

Lippert-Rasmussen, K. (2006) 'Racial Profiling Versus Community', *Journal of Applied Philosophy*: 23 (2), 191-205.

Medina, J. (2011) "The Relevance of Credibility Excess in a Proportional View of Epistemic Injustice: Differential Epistemic Authority and the Social Imaginary', *Social Epistemology*: 25, 15-35.

Moss, S. (2018), 'IX—Moral Encroachment', *Proceedings of the Aristotelian Society*: 118, 177-205.

National Academies of Sciences, Engineering, and Medicine. 2018. "Racial Bias and Disparities in Proactive Policing." *Proactive Policing: Effects on Crime and Communities*. Washington, DC: The National Academies Press.

Pohlhaus, G. (2011) 'Relational Knowing and Epistemic Injustice: Toward a Theory of Willful Hermeneutical Ignorance', *Hypatia:* 27 (4), 715–35.

Singh, K. (2022) ‘There is No Pragmatic Encroachment’, *Unpublished Draft*.
Spear, A. (2019), ‘Epistemic Dimensions of Gaslighting: Peer-Disagreement, Self-Trust, and Epistemic Injustice’, *Inquiry: An Interdisciplinary Journal of Philosophy*: 1-24
Spencer, K. B., Charbonneau, A. K., and Glaser, J. (2016) ‘Implicit Bias and Policing’, *Social and Personality Psychology Compass*, 10: 50– 63.
Toole, B. (2019), ‘From Standpoint Epistemology to Epistemic Oppression’, *Hypatia:* 34(4), 598-618.