

## BEYOND ACCURACY: EPISTEMIC FLAWS WITH STATISTICAL GENERALIZATIONS<sup>i</sup>

Jessie Munton

St John's College, University of Cambridge

What, if anything, is epistemically wrong with accurate statistical generalizations about demographic groups? This paper argues that there is a perfectly general, underappreciated epistemic flaw which affects both ethically charged and uncharged statistical generalizations. Though common to both, this flaw can also explain why demographic statistical generalizations give rise to the concerns they do. To identify this flaw, we need to distinguish between the accuracy and the projectability of statistical beliefs. Statistical beliefs are accompanied by an implicit representation of the statistic's modal profile. Their modal profile determines the circumstances in which they can legitimately be projected to unobserved instances. Errors in that implicit content can be compatible with the accuracy of the "bare" statistic, whilst systematically leading to downstream errors in reasoning, in a manner which reveals an epistemic flaw with an important aspect of the belief state itself.

There are some generalizations about demographic groups which we may hesitate to endorse, even as they are borne out by empirical data. These include claims like the following: "Black Americans are almost eight times more likely to have a homicide conviction than white Americans;"<sup>ii</sup> "teenage girls perform less well at mathematics than boys;"<sup>iii</sup> "gay men have far higher rates of sexually transmitted diseases than straight men;"<sup>iv</sup> "conservative political views are associated with lower IQ."<sup>v</sup>

These sound like classic racist, sexist or homophobic claims, and yet they are statistically supported. Nonetheless, we tend to feel uneasy with such beliefs, and to suspect that someone who drew on them in their reasoning about the people they met would be rationally remiss.<sup>vi</sup> This gives rise to the following puzzle: what explains that unease? Is it due merely to a socially commendable but epistemically unwarranted fear of saying the unsayable, even a regrettable concession to political correctness at the expense of articulating the truth?<sup>vii</sup> If it reflects a genuine flaw with the beliefs themselves, to what can we trace it given their accuracy?

Accurate but intuitively problematic beliefs about demographic groups such as these have been the focus of recent work in philosophy, focusing on their ethical implications. Some have taken our

unease with beliefs such as these as evidence that our beliefs, like our actions, can be the subject of moral evaluation. Merely *holding* a belief about a group of individuals, or drawing inferences about an individual on the basis of such beliefs can morally wrong the individuals concerned (Basu, (forthcoming), Preston-Roedder (2013), Rinard (2015)). They argue that the possibility of such a wrong makes us wary of endorsing these beliefs despite their epistemic good-standing. Relatedly, other recent work has raised the interesting possibility that our response to these beliefs reveals that our epistemic standards are not independent from the ethical significance of the beliefs in question (Basu 2019, forthcoming a; Bollinger 2018). Rather, these beliefs may fall foul of an epistemic standard *in virtue of* their ethical significance. The ethical significance of beliefs about race or gender inflates the epistemic standards in play.<sup>viii</sup> By attributing negative properties to demographic groups, we risk encouraging unfair differential treatment of those groups, or limiting our capacity to treat individuals equally. On this approach, those risks inflate the standards a belief must meet if it is to be justified, or qualify as knowledge.<sup>ix</sup>

By contrast with these approaches, I want to argue that we can locate a relatively straightforward epistemic problem with many such beliefs, even before we consider their moral dimensions or the downstream problems they give rise to. Mere accuracy need not be the end of the road for the epistemic evaluation of these beliefs. To this extent, the account I offer here is compatible with these morally grounded alternatives; but it is also a corrective, in so far as we should resist the rhetorical implication of these alternative approaches that we *can only* appeal to ethical features, because no rational flaws remain to be identified with beliefs of this sort.

I argue that there is widespread scope for an epistemic flaw with these beliefs, *in spite of* their apparent accuracy and rationality, and *in addition to* the moral unease that can appropriately co-exist with those epistemic virtues. To do so, I establish the potential for an epistemic flaw in parallel cases of statistical generalizations that are without moral or political charge, and I show how this epistemic flaw can affect these "problematic" cases too. We reason badly with such statistics not just from a willful

misconstrual, but because of a more basic fact about the nature of statistical beliefs that creates a rational vulnerability, which ideology can then exploit.

The paper proceeds as follows: in section one I describe more precisely the set of beliefs I am interested in. In section two I offer a comparison case that demonstrates how beliefs involving statistical generalizations can be rationally flawed even when there is no ethical charge to their subject matter, reinforcing the case that the flaw with these beliefs may be epistemic rather than ethical. In section three I offer an account of that flaw in terms of a failure of projectability that is compatible with the accuracy of the bare statistic. In section four I show how that applies to cases of statistical generalizations about particular demographics in particular.

## 1. The target beliefs

Here are three examples of the kinds of beliefs I am primarily concerned with in this paper.

*Black homicide:* Nina lives in a largely white community, and seldom interacts with non-white people. She reads in a variety of reliable news sources that 52% of homicide convictions in the United States are of black people, though black people constitute only 13% of the population. On the basis of this evidence, Nina believes that black people in the United States are proportionately more likely to be convicted of homicide than white people.

*Conservative intelligence:* Mike's friend and family are almost exclusively left-wing or liberal. In a political science class at university, Mike learns of various studies that suggest lower IQ is associated with conservative social attitudes, and that right-wingers tend to have a lower level of education than left-wingers. On the basis of this, Mike forms the belief that conservatives tend to be less intelligent and educated than liberals, in line with the statistics he has encountered.

*Abstract thinkers* (adapted from Arpaly (2002)): Solomon lives in a small, isolated farming community. All the women in Solomon's community, despite his attempts to engage them in learned conversation, seem to discuss nothing but gossip, family, and manual work. The community's small, outdated library contains abstract work written by men only. The few people in his community who are interested in abstract thinking are all men, and no one he knows of has ever doubted that women are worse abstract thinkers. On the basis of this evidence, Solomon believes that women in his community are less inclined towards abstract thinking than men.<sup>8</sup>

These cases share three important features. Firstly, the resultant beliefs are accurate. Secondly, the beliefs attribute a negative characteristic to a particular demographic. Thirdly, the subject forms the belief *casually* but *responsibly*. That is, they are not actively reflecting on the reasons for the generalization in question. At the same time, they rely on a reasonable body of evidence from a reasonably reliable source.<sup>9</sup> I will refer to beliefs of this kind as *statistical generalizations*, a term designed to encompass both beliefs that represent precisely the prevalence of a property amongst a population (*74% of x's are y*), and beliefs that do so in more rough, general, or comparative terms (*most x's are y*, *x's are twice as likely to be y as z's*).

Importantly, this set of beliefs goes beyond generic claims about these groups, though it may include some such claims: the problem I am interested in also arises for precise statistical claims, or probabilistic claims about these groups.<sup>10</sup> In addition, discussion of generics has tended to focus on the effects of *asserting* these sentences.<sup>11</sup> I am interested in the epistemic status of beliefs which may never be asserted. That limits the resources we have to locate a flaw in them: our explanation of what goes wrong when individuals form these beliefs cannot be framed in terms of conversational pragmatics, for instance.

Nor does the set of beliefs I am interested overlap exactly with beliefs that might be termed "stereotypes", understood as associations between particular social groups and certain characteristics (Beeghly 2015, Puddifoot 2017). The beliefs I am interested in are a far more precise subset, *statistical*

beliefs about particular demographics. Though these sorts of statistics may lend support to stereotypes, because of their greater precision, they pose a sharper and more focused problem than stereotypes *per se*.

Note, finally, that we are concerned with beliefs about a population or group. None of our thinkers are yet applying these statistics to a particular individual. It may be that there are additional epistemic problems that arise at that stage (see Moss (2018), Bolinger (2018)). I am concerned with our sense that there is something awry *even before the thinker applies the generalization to a particular individual*, though that application to an individual may importantly *reveal* the flaw in question.

## 2. Dalmatians and Chinese Elms

I use the term *epistemic* to indicate a flaw that arises in virtue of an irrational response to information, or through a lack or loss of information. If our unease with these beliefs reflects not just an ethical flaw linked to their distinctive subject matter but a genuinely epistemic flaw, then we should expect it to arise for other statistical generalizations, ones that are unconnected with the sorts of properties which ground discriminatory beliefs and actions. I turn now to two such examples. Diagnosing the problem in these cases will help us identify it in the target set we started from.

*Dalmatians:* According to the Dalmatian club of America “responsible breeders require spaying or neutering of companion animals.... Only the top-quality show-specimens should be considered for breeding.” As a result, in excess of 95% of Dalmatians in the United States are spayed or neutered, leaving them infertile. Chris has a merely passing interest in dogs but a keen nose for statistics. He hears from a reputable source that only 5% of Dalmatians in the US are fertile. He forms the belief on the basis of this report that only 5% of Dalmatians in the US are fertile.<sup>xiv</sup>

*Chinese Elms:* Chinese Elm trees are native to eastern Asia, where they naturally grow to between 15 and 20 meters tall. They are found only rarely in the wild in the USA.<sup>xv</sup> Chinese Elm trees are one of the most popular species to be cultivated as a Bonsai tree, however. Many Chinese Elm trees are in the United States in bonsai form. These bonsai trees rarely exceed 18 inches in height. As a result, 98% of Chinese Elms in the US are under 18 inches in height. Tristan is no dendrologist, but is told this statistic by a friend who is keen on gardening, and comes to believe it on that basis.

Chris and Tristan hold accurate beliefs. The beliefs are reasonably well founded. Neither Chris nor Tristan seems rationally remiss or blameworthy for holding these beliefs.

Nonetheless, there is a rational flaw in the offing, which becomes evident when we imagine Chris and Tristan drawing on their belief in performing certain inferences. Imagine Kohava phones up Tristan, concerned about a small tree sprout she’s found in a corner of her garden. “I think it’s a Chinese Elm,” she tells him. “I hope it isn’t going to grow and block the light.” “I wouldn’t worry!” Tristan tells her. “Chinese Elms rarely grow beyond 18 inches in the States. You’d have to be very unlucky for this one to do so.”

Meanwhile, Glenda phones Chris, asking his advice about a stray Dalmatian she’s brought back to the States from Croatia. “Should I get him spayed?” she asks. “Don’t worry!” Chris tells her. “Very few Dalmatians are fertile. I really doubt it’s worth the money.”

## 3. What went wrong?

What explains the flaw in Chris and Tristan’s reasoning, given that the statistical generalizations they employ are accurate? Dalmatians and Chinese Elms are instances of a familiar problem: they are occasions on which statistical syllogism fails. What features of the statistic that features in the inference predict its failure, if not its inaccuracy?

Chris and Tristan's reasoning brings them up against a problem that I shall call the *Problem of Projectability*. They have each projected a regularity observed among past instances of a kind to a novel, unobserved instance. In doing so, they intuitively project the statistic too broadly, onto an item to which it does not apply. The problem of projectability is the problem of when it is legitimate to project a statistic onto a novel instant that falls under the description used to initially characterize the set in question.

Chris and Tristan's beliefs are accurate. But they are not projectable as they currently stand: when Chris and Tristan try to project their beliefs forward onto unobserved instances, they are liable to do so in ways that lead them astray. To understand the scope for epistemic flaws with statistical generalizations, it is key that we notice that accuracy and projectability can come apart from one another.

The statement of a statistic employs a description of the group in question. That description picks out a set of individuals: that is the domain of the statistic for our present purposes. A statistic is accurate if the appropriate proportion of items in the domain manifest the property in question (perhaps over some relevant time frame). The domain is commonly an *extensional* notion: it picks out a set in the actual world.<sup>xvi</sup> When we come to project the statistic onto a novel, unobserved instance, however, we have to appeal not to an extensional notion, but to an implicit modal interpretation of the statistic. When projecting a statistic, we are bound to assume something about the counterfactual stability of the relevant regularity, including what, if any, further facts underpin it. We can call this its *modal profile*.

The accuracy of a statistic may depend on particular idiosyncratic and contingent features of the domain. Failing to take those contingencies into account when projecting the regularity onto a novel, unobserved instance leaves one liable to overproject. That can be the case even as the initial statistic is accurate. For instance, the fact that a statistic holds of a given domain under some description or other may be largely coincidental, in the sense that there is nothing about the description in question which *causes* or *explains* the statistical regularity. It so happens that 75% of objects currently on my desk have blue on some part of them, but there is no explanation or cause that relevantly links the objects which fall under that description with that particular property.<sup>xvii</sup> Alternatively, the regularity may depend, as it

does in the cases of Chris and Tristan, on a particular intervention to which members of the domain have been subjected, but which is not made explicit in the description which picks out the domain in question.

This kind of coincidence, or dependency on a particular intervention, limits the modal profile of the statistic in question. Just as a scientist has to consider the external validity of their findings, (that is, the extent to which a finding will generalize to different individuals in different circumstances from those tested in the experimental sample,) so we all, when holding or reasoning with statistical beliefs, have to consider how modally robust the generalization in question is.<sup>xviii</sup> The mere accuracy of a statistic does not guarantee that it will have sufficient modal stability that it can be reliably projected onto a given novel instance.

Responsible projection, then, requires some representation of the modal profile of the statistic in question. That involves a raft of further implicit content that has to supplement the bare statistic, and that may include tacit explanatory or causal representations, for instance. When statistics are misprojected, that need not be because the bare statistic the subject endorses is inaccurate. A believer may instead misrepresent the modal profile of the statistic, and over- or under- project it on that basis, in spite of the accuracy of the simple statistic itself.

On this approach, we can distinguish between different elements in a statistical belief. On the one hand, there is the more limited construal of its content as the simple attribution of a property to a set with some or other frequency. That allows us to determine if the statistic is accurate for a particular domain singled out under some description or other in the actual world. But in addition to that explicit content, any statistic is accompanied by a further implicit representation of its generality: how modally robust the statistic is, in virtue of its dependency on particular contingencies. That implicit content is nonetheless an essential part of a statistical belief, in so far as it is required for a subject to reason with or project the belief in question. The former content may be accurate whilst the latter is not. An important upshot of this account is that the most naturally reported description of the relevant belief is really the tip of an iceberg, in the sense that it is a small part of a network of representations which provide

additional implicit, sometimes explanatory, content. The epistemic good-standing of a belief depends on what is going on “under the water”, in that broader belief structure.<sup>318</sup> Simply claiming that a statistic is accurate is insufficient to establish that the subject’s belief representing that statistic is epistemically a-ok. We need to enquire in addition into whether the subject appropriately circumscribes the modal profile of the belief, since that will determine whether or not they reason well with it.

How does this apply to the cases of Chris and Tristan? Their statistical beliefs are accurate, when assessed over the relevant domain in the actual world. However, the broader web of implicit content which supplies the tacit representation of the statistics’ modal profile contains inaccuracies. It is on the basis of that tacit representation that they assimilate or excludes novel dogs or trees from the umbrella of the original statistic. Hence, errors in that tacit representation explain their misprojection of the statistic in question. Chris and Tristan’s overestimation of the statistic’s modal strength constitutes an epistemic flaw in their belief, even as it does not impugn the accuracy of the bare statistic.

Chris and Tristan’s assumptions are not uncommon. Statistics are often taken not just to hold over a sample of kind distinguished by the fairly arbitrary parameters of time and place, but in virtue of some more fundamental commonality that licenses the forward projection of the statistic onto other instances. In fact, the domain of a particular regularity is frequently circumscribed by social practices that shape the category to which the property in question is attributed. Dalmatians and Chinese Elms are subject to particular social practices. Those practices underwrite the truth of many generalizations about them. They determine facts about their size, health and fertility, for instance. This point isn’t limited to members of the animal or plant world. True statistics about street lamps and body parts, houses and food all depend on certain social structures and practices obtaining. These structures and practices circumscribe the modal profile of the resulting statistics: in nearby worlds in which the relevant social structures and practices fail to obtain the regularity no longer holds.

We can therefore trace the erroneous beliefs Chris and Tristan arrive at concerning the individual instance of a kind – that it will not grow above a certain height, or will very probably be infertile—to more than the kind of bad luck that is an ever-present possibility when reasoning with statistics. The inaccurate result was a *predictable* upshot of a flaw higher up the chain of reasoning: they over-estimated the modal strength of the statistics, and *that* is what gives rise to their inappropriate application. That implicit inaccuracy leads to systematic and predictable errors in inferences about unobserved instances.

#### 4. Extending the analysis to demographic statistics

How does this help us with the cases of demographic statistical generalizations described in section one? Recognizing that accuracy and projectability can come apart opens up room to acknowledge that the accuracy of statistical beliefs about particular demographics is compatible with systematic failures of projectability, when accuracy over some domain is coupled with an inaccurate implicit representation of the modal profile of the statistic in question. An implicit awareness of that potential disconnect between accuracy and projectability fuels our concern with sanctioning even accurate statistics about particular demographics. Endorsing the statistic in question without making explicit the limitations of its modal profile can imply that it, and inferences from it, enjoy a stronger epistemic standing than they really do.

Given the breadth of this potential flaw, we still need to understand why this should afflict statistics about demographic groups in particular. The scope for this separation between accuracy and projectability is particularly acute in these cases because such statistics often depend on contingent features of social organization. That contingency limits the modal profile of the statistic, and enlarges the scope for divergence between statistical accuracy and projectability. At the same time, we tend to essentialize observed differences between demographics, failing to do justice to their dependence on particular social structures.<sup>319</sup>

Take Solomon's belief that women in his community are not inclined to abstract thinking. His belief is true only in virtue of a particular set of facts about the social structure in which he and they are located, which deprives them of the opportunities required to develop such an inclination. Those facts limit the modal profile of the statistic, and with that its projectability. If Solomon is to appropriately reason with the statistic, he must represent, at least implicitly, the dependency of the regularity on that particular set of social structures and practices. Otherwise he is likely to expect the generalization to hold outside of the relevant social structures, and hence to draw erroneous conclusions from it.

In one sense this point should be familiar. Via its entwinement with explanatory interpretations of certain statistics, it is a central theme in the long history of scholars writing on racial prejudice. Khalil Muhammad repeatedly demonstrates this much in his book-length discussion of the roots and consequences of the association between blackness and criminality in modern American. Muhammad (2011) painstakingly chronicles the ways in which early discussion of black criminality tended to attribute it to innate racial traits, whilst discussants of white criminality were ready to attribute it to circumstantial factors such as social conditions. Frederick L. Hoffman's 1896 *Race Traits* was a case in point. Hoffman interpreted self-destructive behavior, high disease and mortality rates amongst the black population as evidence of genetic racial inferiority, whilst blaming social and economic factors for similar phenomena among white populations. W.E.B. DuBois countered in *The Philadelphia Negro* (1899) with a class-based analysis that allowed him to present black criminality as the result of economic conditions rather than racial factors.

This disagreement makes explicit how agreement about the accuracy of the statistic in question is compatible with disagreement in this case about explicit representations of its modal profile and consequent projectability, in virtue of disagreement about the appropriate explanation for the statistic in question. As the debate evolved throughout the 20<sup>th</sup> century, what Muhammad calls the "statistical reality of black criminality" (2011: 68) was seldom at issue. Rather, it was the reasons *for* those statistics that were contested. In so far as some representation of generality or modal stability is an intrinsic part of the

belief state, the disagreement between DuBois and Hoffman should be understood not just as a disagreement about extrinsic material, but as a disagreement about the nature of the statistic itself, even as that is compatible with some agreement about the bare incidence of the trait in question in the relevant domain.

This disagreement between DuBois and Hoffman highlights another relevant feature of statistics about demographic groups. Often the modal profile of a statistic is limited not just by facts about social structure *per se* but a specifically *unjust* social structure. That unjust social structure systematically slants the opportunities available to individuals within a society away from some groups and towards others. This gives rise to particular epistemic problems: suppose someone is raised in a society replete with structural injustice towards a certain population. This may lead to a disposition to hypothesize a common cause for multiple negative statistics about that population. Under these conditions, one might be led to think that the common cause is something about the *nature* or *essence* of the population in question. Essential features of a population hold regardless of context. So, if one is living in a society replete with structural injustice towards a population, one will be disposed to take a negative statistical fact about that population to be evidence for a claim that holds with much greater generality.

But there is nothing unique about the epistemic impact of structural injustice in this respect. The point is perfectly general. To reason well on the basis of certain statistical generalizations, your reasoning needs to be constrained by an appropriate (implicit) representation of the modal profile of that statistic. Those facts determine the inferences it is appropriate to draw from the statistic. That is what makes this particular problem with certain demographic generalizations a genuinely *epistemic* problem: it resides in a lack of information, not in the morally valenced subject matter of these generalizations. The lack of information leads to the misconstrual of the statistic's modal profile, and a flawed pattern of inferences based on that inaccuracy. On the other hand, a failure to represent the role of structural social injustice in bringing about certain facts is liable to further perpetuate that injustice. This is one way in which these cases *are* particularly ethically charged. Our unease is intensified by the potential for this inaccuracy to be

mobilized to support prejudicial beliefs and behavior. We are in a better position to describe those downstream ethical flaws when we appreciate that they are grounded in a more basic informational deficit. That ethical charge makes us particularly sensitive to the epistemic flaw in these cases.

Importantly, the problem is not limited to the possibility that one draw inappropriate inferences about an individual's *essence*. We can equally go astray in construing a statistic's modal profile by attributing the prevalence of a property to the *wrong* non-essential cause. Tristan will reason no better about the appropriate treatment of his friend's pet if he thinks that Dalmatians spay themselves, or that the neighborhood cats will perform the intervention. More pertinently, the Christian right, for instance, has tended to attribute negative statistics to members of the LGBTQ community, whilst insisting those are the result *not* of an individual's homosexual essence (since they deny anyone is homosexual by nature) but of a broader cultural of immorality.<sup>xvi</sup> Explanations, implicit or explicit, that root a generalization in culture rather than essence are just as capable of giving rise to errors in the construal of the modal profile of the statistic.<sup>xvii</sup>

This approach is also capable of teasing out why some demographic generalizations strike us as more problematic than others. Women are on average shorter than men. That statistic is more modally robust than statistics around women's academic performance. It is less dependent on the social structure in which they are embedded. As a consequence, it does not strike us as potentially problematic in the same way as the generalization that women don't code, or do mathematics at the highest levels, a generalization which plausibly relies at least in part on unjust social structures, and whose modal profile is more limited as a result. Projection on the basis of the first statistic is less vulnerable to error than projection on the basis of the second, though both may be accurate of women in the actual world.

This also helps explain why whether a statistic seems epistemically problematic depends on the context in which it is deployed, and the identity of the believer. We could imagine a member of the Black community researching for a non-profit intervention endorsing the same statistic about race and criminality as Nina whilst raising fewer qualms. Nina forms her belief about race and crime casually. She

does not attend to its modal profile. She is vulnerable then to misrepresenting it, and drawing incorrect inferences from the belief as a result. The community member and researcher, on the other hand, is more likely to actively attend to the facts in virtue of which the statistic in question holds and to appropriately represent its modal profile. As a consequence, the scope for her belief to be epistemically flawed is more circumscribed.

Note that that does not render it impermissible to hold these beliefs, it only requires that we take steps to accurately represent limitations on the generality of the statistics. That is why these problem cases tend to involve beliefs that are formed casually: the risk of inaccuracy in the additional material drawn on to determine the statistic's modal profile is lessened when that generality is explicitly at issue. The problem arises most obviously in contexts in which these beliefs are held without any consideration of its generality.

## Conclusion

The account offered here identifies a new dimension along which we can epistemically evaluate statistical beliefs. To more standard axes of evaluation—the accuracy or reliability of the bare statistic, we can now add projectability. But this raises the question: how does projectability impact on the justificatory status of statistical beliefs? Can a statistical belief which is accurate but not projectible be doxastically justified? On most standard approaches, doxastic justification depends on the *explicit* content of a statistical belief, whereas the representation of modality which determines its projectability is frequently *implicit*. Moreover, it is tempting to think of a failure of projectability as a flaw not with the belief itself, but with the inferences drawn from it. Allowing what happens *downstream* from a belief to decide whether the belief itself is justified looks like a departure from standard approaches to justification on which it is the *genealogy* of the belief which primarily determines its epistemic status.

I've argued, however, that those flawed inferences predictably result from problems with the belief state itself, namely, tacit inaccuracies in the implicit representation of its modal profile. And that

implicit representation is an unavoidable component of any statistical belief. When beliefs are not projectable, there is a tacit inaccuracy in the belief state. Given the intimate and essential nature of that implicit modal representation, a failure of projectability reflects an epistemic flaw with the belief itself, even as it may coexist with the superficial accuracy of the bare statistic. It therefore appropriately detracts from the epistemic status of a belief as justified or unjustified. Justification should be sensitive not just to the top of the iceberg – the explicit belief content – but to the bottom of the iceberg too – the mass of tacit content that determines the modal profile attributed to the belief.

Recognizing that projectability legitimately features in our evaluation of our beliefs lets us recognize the intuitive unease we feel with these beliefs as a response to an important feature of statistical generalizations, their modal profiles. The unease is a sensitivity not just to the potential for hurt feelings but to a crucial aspect of proper statistical belief formation. These beliefs are not fully justified, despite their superficial accuracy.

Of greater import, however, is not how we apply the term justification, but what follows from this account about the nature and extent of our epistemic obligations. The crucial upshot is that we are under an obligation to ensure the accuracy of our tacit representation of the a statistic's modal profile, at least to a level that renders the inferences we draw on the basis of that belief reasonably reliable. Doing so requires a sensitivity to whatever underpinning conditions secure the boundaries of its modal profile. Otherwise, a subject is liable to reason as though the statistic reveals, for instance, properties intrinsic to the group of which it holds true, when it in fact arises only as the result of contingent features of social structures or organization.

The facts which determine a statistic's modal profile are frequently highly complex. An epistemic obligation to track that complexity can seem a high bar to meet. This complexity may seem to make the claim there is any such obligation implausible. If such an obligation existed, wouldn't we be widely in breach of it?

Our epistemic obligations may well be complex, and we may frequently fall short of them. That fact alone does not get us off the hook. It may, however, alter the extent to which we regard this as a reason to blame an individual if they fall short of it unwittingly. The epistemic obligation may lie instead, at least in part, on the social structures on which the individual reasonably relies to provide them with relevant information. But questions about the proper distribution of culpability for failures of this kind do not change the fact that the epistemic good standing of statistical beliefs depends on the generality attributed to the statistic in question.

This analysis helpfully allows us to pinpoint a range of ways in which beliefs endorsing statistical generalizations can be epistemically flawed. But it is unlikely to be exhaustive. There will be further epistemic flaws that arise when we apply a statistical generalization to an individual. There may be a host of additional ethical flaws that are responsible in turn for, for instance, flawed evidence gathering or processing, or inappropriate action on the basis of true beliefs. It is possible too that ethical stakes raise epistemic standards. These solutions compliment rather than compete with one another. The flaw described here is an important piece of a complex puzzle.

- Arpaly, N. (2002). *Unprincipled Virtue: An Inquiry into Moral Agency* OUP
- Basu, R. (forthcoming) The Wrongs of Racist Beliefs. *Philosophical Studies*
- Basu, R. (forthcoming a.) Radical Moral Encroachment: The Moral Stakes of Racist Beliefs. *Philosophical Issues*.
- Basu, R. (2019) Doxastic Wronging. In Kim, Brian and McGrath, Matthew (eds.) *Pragmatic Encroachment in Epistemology*. Routledge. pp.181-205
- Beehly, E. (2015). What is a Stereotype? What is Stereotyping? *Hypatia* 30.4 675-691
- Bolinger, R. (2018). The Rational Impermissibility of Accepting (some) Racial Generalizations. *Synthese*.
- Bortolotti, L. (2015). The Epistemic Innocence of Motivated Delusions. *Consciousness and Cognition* 33: 490-499
- CDC Centers for Disease Control and Prevention (2016, March 9) Gay and Bisexual Men's Health: Sexually Transmitted Diseases. Available at <https://www.cdc.gov/msmhealth/std.htm>
- Cooper, A. and Smith, E.L. (2011). "Homicide Trends in the United States, 1980-2008" *Bureau of Justice Statistics* U.S. Department of Justice NCJ 236018. Available at <https://www.bjs.gov/content/pub/pdf/htus8008.pdf>.



Dotson, K. (2014). Conceptualizing Epistemic Oppression. *Social Epistemology* 28 (2):115-138.

Du Bois, W. E. B., & Eaton, I. (1899). *The Philadelphia Negro: a social study* (No. 14). Published for the University.

Fricker, M. (2007). *Epistemic Injustice: Power and the Ethics of Knowing* (Vol. 69): Oxford University Press.

Fritz, J. (2017). Pragmatic encroachment and moral encroachment. *Pacific Philosophical Quarterly*. 98 (S1):643-661

Gardiner, G. (2018). Evidentialism and Moral Encroachment. In McCain (ed.), *Believing in accordance with the Evidence: New Essays on Evidentialism*. Springer. pp.169-195

Gendler, T. (2011). On the epistemic costs of implicit bias. *Philosophical Studies* 156 (1):33-63.

Goodman, N. (1955). *Fact, Fiction and Forecast*. Harvard University Press

Hacking, I. (1995). The looping effects of human kinds. *Causal cognition: A multidisciplinary debate*, 351-394.

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

Hájek, A. (2007). The reference class problem is your problem too. *Synthese*, 156(3), 563-585.

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

Hodson, G., & Busseri, M. A. (2012). Bright minds and dark attitudes lower cognitive ability predicts greater prejudice through right-wing ideology and low intergroup contact. *Psychological Science*, 23(2), 187-195.

Hoffman, F. L. (1896). *Race traits and tendencies of the American Negro* (Vol. 11, No. 1-3). American economic association.

Hudson, J. (2015, November 28) “5 Devastating Facts About Black-on-Black Crime” *Breitbart* Retrieved from <http://www.breitbart.com/big-government/2015/11/28/5-devastating-facts-black-black-crime/>

Hume, D. (2003). *A Treatise of Human Nature* New York: Dover

Leslie, S.-J. (2013). Generics Oversimplified. *Noûs* 47 (3):28-54.

Molden, D. C. & Higgins, E. T. (2005). Motivated thinking. In K. Holyoak & B. Morrison (eds.), *The Cambridge Handbook of Thinking and Reasoning*. Cambridge University Press 295--317.

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

Muhammad, K. G. (2011). *The condemnation of blackness*. Harvard University Press.

Nickel, B. (2016). *Between Logic and the World: An Integrated Theory of Generics*. Oxford University Press UK.

Niederle, M. and Vesterlund, L. (2010). Explaining the Gender Gap in Math Test Scores: The Role of Competition. *Journal of Economic Perspectives* 24:2 129-144

OECD. (2014). *Pisa 2012 Results in Focus: What 15-Year-Olds Know and What They Can Do with What They Know*. Available online at <https://www.oecd.org/pisa/keyfindings/pisa-2012-results-overview.pdf>

Perry, M. (2014) The 2014 math SAT test results confirm a pattern that has persisted for 40+ years: boys are better at math than girls. *AEI Publication* Accessed at <http://www.aei.org/publication/2014-math-sat-test-results-confirm-pattern-persisted-40-years-boys-better-math-girls/>

Pinker, S. (2007). In Defense of Dangerous Ideas. *Chicago Sun-Times*.

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

Puddifoot, K. (2017). Stereotyping: the multifactorial view *Philosophical Topics* 45.1:137-156

Puddifoot, K. (2017a). Dissolving the epistemic / ethical dilemma over implicit bias. *Philosophical Explorations* 20 sup1: 73-93

Rinard, S. (2015). No Exception for Belief. *Philosophy and Phenomenological Research* 94(1): 121-143.

Saul, J. (2017) Are Generics Especially Pernicious? *Inquiry*

Shelby, T. (2003). Ideology, racism, and critical social theory. *Philosophical Forum* 34 (2):153–188.

Soames, S. (1986.) Incomplete definite descriptions. *Notre Dame Journal of Formal Logic* 349 – 75.

Stanley, J. (2015) *How Propaganda Works*. Princeton University Press

Stanley, J., & Szabó, Z. G. (2000). On Quantifier Domain Restriction. *Mind and Language*, 15(2&3), 219–261.

Waugh, R. (2016, 31 January) Right-wingers are less intelligent than left winger, says study. *Daily Mail* Available online at: <http://www.dailymail.co.uk/sciencetech/article-2095549/Right-wingers-intelligent-left-wingers-says-controversial-study--conservative-politics-lead-people-racist.html>

<sup>i</sup> For feedback on earlier versions of this paper I am indebted to Jason Stanley, Susanna Siegel, Daniel Greco, Bernard Nickel, Briana Toole, Georgi Gardiner, Rima Basu, Michael Strevens, Lisa Miracchi, Michaela McSweeney, an audience at the 2018 Pacific APA, and participants in the Gdansk LoPSE seminar and the Yale philosophy work in progress seminar.

<sup>ii</sup> Hudson (2015). Blacks were convicted of 52% of homicides between 1980 and 2008 in the US, whilst constituting 13% of the population (Cooper and Smith 2011).

<sup>iii</sup> Perry (2014). In OECD countries, they underperform their male counterparts by an average of 11 points (OECD 2014:23). The ratio of boys to girls in the top 5% of math performance is two to one (Niederle and Vesterlund 2010)

<sup>iv</sup> Gay men accounted for 83% of syphilis cases in the US in 2014 and are 17 times more likely to suffer from anal cancer than straight men (CDC 2016).

<sup>v</sup> Waugh (2016). See Hodson and Busseri 2012 for evidence that conservative ideology accounts for the predictive relationship between lower IQ in childhood and racist attitudes in adulthood.

<sup>vi</sup> Of course, we are idealizing away from a host of common or garden rational flaws which routinely, albeit contingently, afflict many actual instances of beliefs with this kind of content: we are bad at reasoning with and often overproject such statistics (Puddifoot 2017). Beliefs such as these often arise from flawed ideological processes of belief formation, not as a proportionate response to evidence.

<sup>vii</sup> Steven Pinker, for instance, suggests that certain “dangerous ideas”, including the suggestion that women have, on average, a different profile of aptitudes and emotions, and that African American men have higher levels of testosterone, are “denounced not because they are self-evidently false, nor because they advocate harmful action, but because they are thought to corrode the prevailing moral order” (2007).

<sup>viii</sup> See also Moss (2018) for a version of this response. For an argument in favor of moral encroachment more broadly see Fritz (2017). For discussion and criticism of arguments in favor of moral encroachment see Gardiner (2018).

<sup>ix</sup> Another strand of work locates our intuitive unease in the downstream problems that beliefs such as these can give rise to. Tamar Gendler (2011), for instance, argues that reliance on base rates and racial categories has significant epistemic costs, even as it is mandated by our cognitive finitude. Such reliance results in our losing information, in virtue of features of our psychology and our resulting vulnerability to stereotype threat, for instance, the same race face effect, or cognitive depletion brought on by the effort of repressing negative content associated with stereotypes. Katherine Puddifoot (2017a) offers an alternative analysis of beliefs that endorse stereotypes, arguing that the downstream costs are so high that failing to encode base rates may be “epistemically innocent”, in Lisa Bortolotti’s (2015) terminology, and that the best ethical option consequently aligns with the best epistemic option in these cases. These kinds of beliefs may also make it harder to gather information in the future by giving rise to forms of epistemic oppression (Dotson 2014) and injustice (Fricker 2007) that in turn perpetuate epistemic and ethical flaws.

<sup>x</sup> Importantly, Nomy Arpaly sets up the case so that Solomon’s belief is false of the population in his environment. I am interested in whether it may be epistemically flawed even when it is true of some relevant sample. See Miranda Fricker (2007) 32ff for further discussion of the case as presented in Arpaly (2002).

<sup>xi</sup> *Ex hypothesi*, these beliefs are accurate and capable of constituting an appropriate response to the evidence at the individual’s disposal. This is certainly not intended to imply, however, that racist or sexist or homophobic beliefs are in general a proportionate response to the evidence. It is hard to isolate our judgments about these cases, in which we suppose the subject has responded appropriately to a body of evidence, from the more common real-world scenario in which individuals who hold such beliefs do so as the result of some kind of irrationality. Beliefs with similar content may often as a matter of fact possess the obvious epistemic flaw of being resistant to evidence, driven instead by deep-seated ideological commitments or

motivated reasoning. For an overview into some of the psychological literature on this see Molden and Higgins (2005). For a philosophical perspective see Puddifoot (2017) Shelby (2003) or Stanley (2015).

<sup>xii</sup> For a sample of the philosophical literature on the problems posed by generics see in particular Leslie (2013), Haslanger (2012) and Nickel (2016). See Saul (2017) for discussion of the fact that the supposed problem of generics affects more than just generic claims.

<sup>xiii</sup> See Haslanger's (2012) account of the problem of certain generic claims in terms of the essentialising pragmatic content associated with assertions of these terms. Haslanger's discussion remains pertinent and I shall return to it in section three below.

<sup>xiv</sup> This is an estimate based on the facts that Dalmatians typically have litters of between 9 and 13 puppies, and female Dalmatians can be bred every other year. Very few Dalmatians are required to be fertile for the Dalmatian population to remain stable.

<sup>xv</sup> This is not true, but we can harmlessly suppose it to be so for the purposes of this paper.

<sup>xvi</sup> The notion of domain as I use it here bears comparison to its use in the context of quantification. A domain of quantification attaches to a sentence, whereas a domain in the sense I use it here attaches to the propositional contents of a belief, where those beliefs involve statistical generalizations, rather than quantifying expressions. Stanley and Szabo (2000) argue that quantifier domains are best understood as properties (and see Soames (1986) for related considerations), whereas I take the domain of a statistic to be a set of individuals. This last distinction may not be a deep one however, since the union of sets a property determines will itself be a set of possible individuals.

<sup>xvii</sup> This example is loosely taken from Nelson Goodman's 1955 discussion of the new riddle of induction. What is the relationship between the problem of projectability and the problem of induction? The problem of induction asks how we know that, as Hume puts it, "instances of which we have had no experience must resemble those of which we have had experience" (2003:1.III.VI). Projectability is familiar from Goodman as a response to the problem of induction: projectable predicates are ones over which it is legitimate to perform inductive inference. The cases we are concerned with here draw attention to the way in which projectability is a *modal* feature, in contrast to accuracy, which need not be.

<sup>xviii</sup> The problem of projectability for belief resembles the reference class problem for frequentist accounts of probability. The reference class problem is the problem of assigning a probability to an event which admits of multiple classifications, when the classification we adopt will in turn affect the resultant probability (Hájek 2007). I am suggesting that a version of this problem arises whenever we observe a regularity among a population under some description: the domain of a statistic admits of multiple competing construals, and which we adopt will determine which inductive extensions of the statistic are ostensibly legitimate and which are not.

<sup>xix</sup> In this respect, the account of the flaw offered here is closely allied to Nickel's 2016 account of generics. Nickel writes that "[a] generic is the tip of an iceberg. It rests on a large theoretical structure that provides the explanatory resources to distinguish the normal from the abnormal, all the while floating beneath the surface of our conscious awareness" (2016:244). My account echoes this claim that apparently simple beliefs draw on more complex explanatory structures to derive their content. Rather than applying this claim to generics in particular, I am arguing that entertaining even a simple statistical belief draws on a rich web of further belief and behavior. The avowed statistical belief is also "the tip of an iceberg".

<sup>xx</sup> The causal chains that determine the generality of these statistics are also frequently looped: a belief that women are less able at mathematics could lead to behaviors that further restrict girls' access to educational opportunities that would improve their mathematical ability. By limiting their mathematical skills, that behavior in turn provides further grounds for the belief itself. This kind of circularity is particularly problematic when the regularity so described is not a mere fact about the natural world, but a fact about the social world that is in part in turn constructed by the perceptions and beliefs of its participants. Haslanger draws on Ian Hacking's work (1995, 1999) on the "looping effects" of social kinds and structures, that is, the way they are self-reinforcing, when analyzing generics. She writes that "...they describe the world as if it is, by its nature, how we have interpreted it, and from there caused it, to be.... In purporting just to capture the facts, the generics import an explanation, implicate that the source of the truth of these claims lies in what cows, women and blacks, *are*" (2012: 468). This is an equal problem for more precise statistical generalizations.

<sup>xxi</sup> For an illustrative example of such a claim, see the blogpost "15 reasons why homosexuality is wrong and hurts society" <http://unsettledchristianity.com/15-reasons-why-homosexuality-is-wrong-and-hurts-society/>. Reason six is that "homosexuals have a much higher incidence of domestic violence", whilst reason 14 is that "homosexuality is not genetic". The author writes that "those who affirm that it is their nature or that they were born that way are only admitting that they are controlled by their sin nature and do not know how to be freed."

<sup>xxii</sup> Relatedly, explanations in terms of culture are just as capable of grounding racist ideology as those couched in terms of essence. See, for instance, Khalil Muhammad's description of the shift that occurred in the early part of the 20<sup>th</sup> century, from attributing black criminality to genetics to "writing crime into culture" (2011: chapter 3).