The Specter of Normative Conflict: Does Fairness Require Inaccuracy?

To appear in *An Introduction to Implicit Bias: Knowledge, Justice, and the Social Mind*, eds. Erin Beeghly and Alex Madva.

Rima Basu
Claremont McKenna College

> Pierce Hawthorne: We all know what we're really thinking. If, and I mean 'if' the culprit is among us, statistically speaking it's Troy.
>
> — Community, *Cooperative Calligraphy*

I. Setting up a conflict

Let's start with an obvious point: we form beliefs about other people all the time. I believe that at a busy intersection no less than three drivers will turn left when the light turns red. Why? Because I see it happen all the time. On similar grounds I believe that when it rains half of my students won't show up for class. Why? Because in my experience no one in Southern California, and I include myself in this generalization, knows how to deal with rain. We often don't think twice about forming beliefs on the basis of these sorts of statistical regularities or stereotypes. But maybe we should.

Consider, for example, whether I should believe that a black man standing outside of a restaurant is a valet. What if it's the case that every valet I have interacted with outside of this restaurant has been black? What if I'm in a rush and I just need to get my car and I know that 90% of valets at this restaurant are black? Although this seems like a classic case of racial profiling, the evidence seems really strong. As Barack Obama has noted in an interview with *People* magazine, "There's no black male [his] age, who's a professional, who hasn't come out of a restaurant and is waiting for their car and somebody didn't hand them their car keys." I might get it wrong, but given the probabilities, I'm also very likely to get it right. But here's the challenge: unjust social structures of our world gerrymander the regularities and the evidence an individual is exposed to in ways that reinforce racist and sexist beliefs (see Munton 2017 and Basu 2018). What if it is the case that because of

historical patterns of discrimination the vast majority of valets working at this restaurant are black? In such a context, if I mistake Barack Obama for a valet, have I done anything wrong?

In one sense, yes. My belief was inaccurate. Further, my belief was an instance of racial profiling, and racial profiling is generally considered to be morally problematic. However, in another sense, in believing as I did I was at least aiming at having an accurate belief. To introduce a technical term, it was *epistemically rational* for me to believe what I did. Suppose, for example, that you check the weather forecast and see that there's a 90% chance of rain tomorrow. Given that 90% chance, you should believe that it will rain tomorrow. Further, it is appropriate to criticize you for being *epistemically irrational* were you to not believe it'll rain tomorrow, e.g., if you started to plan a picnic. By disregarding the likelihood of rain, you are doing something epistemically wrong. It is epistemically irrational and irresponsible to ignore evidence that bears on the question of what to believe.

If we return to the racial profiling case, however, we now find ourselves in the following bind. Believing in accordance with the evidence—the high likelihood that a black man standing outside the restaurant is a valet—may lead you to hold a belief that is morally problematic. After all, assuming that a black man is a valet is a paradigmatic example of a racist belief, and racist beliefs are morally impermissible. The world we live in is pretty racist, and so it shouldn't be surprising if the world presents us with a lot of evidence for pretty racist beliefs. Thus, the stage is now set for the conflict between accuracy and fairness that is the subject of this chapter. Any black male standing outside a restaurant who is mistaken for a valet seems to have the legitimate complaint, no matter how much evidence you had for your belief, that it is *unfair* that you assume that they must be a valet. But now we must ask, why is it unfair?

It's at least partly unfair for reasons that Lawrence Blum has noted. Blum says, "being seen as an individual is an important form of acknowledgment of persons, failure of such acknowledgement is a moral fault and constitutes a bad of all stereotyping" (2004, p. 282, see also Lippert-Rasmussen 2011). Similarly, it's also unfair because we wish to be related to as we are, not as we are expected to be on the basis of race, gender, sexual orientation, class, etc. When someone forms beliefs about us in the same way we form beliefs about planets—that is, as objects to be observed and predicted—they fail to relate to us as persons (see Basu 2019). If you've ever been mistaken for waitstaff or "the help" because, as Pusha T says, your melanin's got a tint, you recognize that feeling of unfairness. If anyone's ever thought that you're not interested in science because you're a woman or made you feel self-conscious

about being sad and wanting to cry because you're a man ("men don't cry"), you know what it feels like to have assumptions made about you. Sometimes they're innocuous or trivial and you don't feel hurt or restricted by them, but often they're not so trivial. Sometimes they really hurt. Although it is notoriously tricky to pin down exactly what is meant by this requirement to treat others as individuals and what precisely it is that makes stereotyping morally wrong (for more see Beeghly 2015 and 2018), we recognize this feeling of being wronged, of being reduced, of being treated as an object, when we are stereotyped. It's, simply put, not fair.

This unfairness of believing something of another person simply on the basis of statistical evidence is also recognized by our judicial system. You cannot convict someone solely on the basis of statistical evidence *no matter how strong that statistical evidence*. Within legal scholarship this is known as the problem of naked statistical evidence (see Schauer 2006, Enoch et al. 2012, and Buchak 2014). The common example goes as follows: while you were driving a bus sideswiped your car. In your town there are only two bus companies: the Blue Bus Company and the Green Bus Company. You want compensation for the damages, so you decide to sue the company operating the bus that sideswiped your car. Here's the problem, though. It was late at night and you couldn't accurately identify the color of the bus. But, you also know that 80% of the buses in the city are blue buses operated by the Blue Bus Company, whereas the other 20% are green and operated by the Green Bus Company. Given the balance of the probabilities, you can be fairly confident that you were sideswiped by a blue bus.

In a civil court you only need to demonstrate a preponderance of evidence. That is, that it is more likely than not that you were hit by a blue bus. The good news is that the statistics are on your side! It is, after all, 80% likely that you were hit by a blue bus and 80% is greater than 50%. The bad news, however, is that this merely statistical evidence is inadmissible in the courtroom. If there had been an eyewitness who could testify that they saw a blue bus sideswipe your car, then you'd stand a much better chance at winning your case *even if that eyewitness testimony was less reliable than the mere statistical evidence.* Why? Again, the answer seems to return to these considerations of fairness. It would be unfair to convict the Blue Bus Company just because it's statistically likely that if a car is sideswiped by a bus it would be sideswiped by one of their buses. It is unfair to hand your keys to a black man assuming that he's a valet just because it's statistically likely that he's a valet.

Although you can't find the Blue Bus Company guilty for sideswiping your car, perhaps you still have enough evidence to believe that it was a blue bus that sideswiped your car. It might be unfair to the Blue Bus Company, but that's where the preponderance of evidence lies. Similarly, the world isn't a fair or just place; the world has been shaped by racism and other discriminatory practices. Perhaps there's some truth to the stereotypes we hold and so if we want to do what's *epistemically right*, i.e., believe in accordance with the evidence, we just have to make this tradeoff with moral considerations like fairness. And this brings us right back to the conflict of this chapter: in the interests of accuracy, must we give up on fairness or in the interests of fairness must we give up on accuracy?

To answer these questions concerning the seeming conflict between accuracy and fairness, we first need to get clear on what is even meant by the claim that you should believe in accordance with the evidence. A common person to turn to when it comes to explicating the duty to believe responsibly is W. K. Clifford. Clifford (1877) asks us to consider an old-timey ship owner whose ship is about to ferry families to their new home. The ship owner knows that the ship is old, it has a few faults and it has been in need of a number of repairs. He's thought to himself that he should have the ship overhauled and refitted, but that would be a costly venture. As he watches the ship leave the dock, he pushes his doubts aside by reminding himself that she's still a hardy ship and she's safely made and returned from a number of journeys. So, perhaps he has no reason to be worried that this time she might not safely come home. I'll leave the last bit of the example in Clifford's own words:

> In such ways he acquired a sincere and comfortable conviction that his vessel was thoroughly safe and seaworthy; he watched her departure with a light heart, and benevolent wishes for the success of the exiles in their strange new home that was to be; and he got his insurance-money when she went down in mid-ocean and told no tales. (p. 1)

Now, is the ship owner guilty of the deaths of the passengers on his ship? He did, after all, *sincerely believe* that his ship was seaworthy. Nonetheless, Clifford argues, the ship owner's sincerity means absolutely nothing because he had no right to believe on the basis of the evidence he had (for more cases like this, see Siegel, this volume). Instead of gathering evidence about the seaworthiness of his ship he instead came to his belief by stifling his doubts. Furthermore, even if his luck (and the luck of the passengers) had been different and the ship had safely made the journey, his guilt would not be diminished one bit. Whether the belief turned out to be true or false has no bearing on whether the ship owner had a right to

believe on the evidence that was before him. As Clifford famously remarks, "it is wrong always, everywhere, and for anyone, to believe anything upon insufficient evidence."

An important part, then, of having epistemically rational beliefs is ensuring that your beliefs are held on the basis of sufficient evidence. It seems pretty intuitive to say that you should believe in accordance with the evidence, and we seem to hold people to this standard all the time. For example, suppose someone were to insist that there has been no foreign interference in our elections. We wouldn't just take their word for it. What we'd want to know is what reasons they have for thinking that, i.e., what evidence they are basing their beliefs upon. But, to now practice some of the pedantry well-known to philosophy: what is evidence? Or to borrow a popular meme, we can present the question as follows.

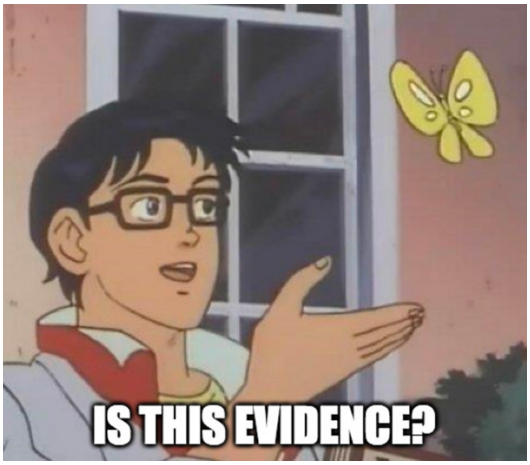

Fig 1.

This is a difficult question to answer, but for our purposes we can appeal to a standard intuitive conception of evidence and how it relates to the *epistemic rationality of belief*. Ordinarily when we think of evidence we think of examples from crime procedurals, e.g., gun residue, fingerprints, DNA from a strand of hair. These examples of evidence are physical objects we can put in a bag, label as evidence, send to a lab, and present in front of a judge or jury. But that's not all that evidence is. What I hear, see, and smell can be evidence that bears on a question under investigation. This question of how to understand the nature of evidence is a big one (see Kelly 2016 for a survey of the topic), but in what follows we will work with the following intuitive gloss: evidence for a question under investigation is a reliable sign, symptom, or mark for information relevant to the question under investigation. For example, smoke is a reliable sign of fire and as such it is *evidence* of fire. Similarly, my dog's barking is evidence of someone at the door, a distinctive whistle from the kettle is

evidence that the water has boiled, and the sight of rain outside my window is evidence that it is raining outside. Now we can fill out our meme as follows.

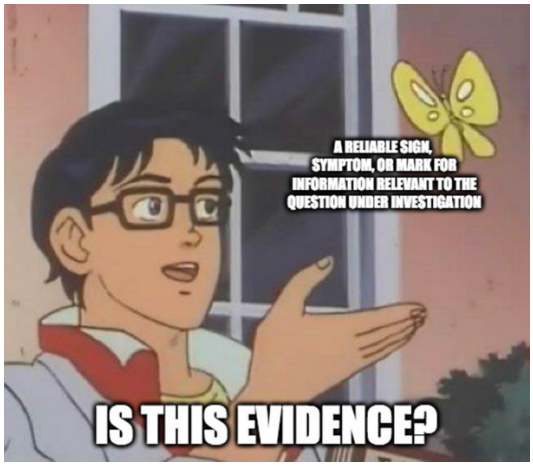

Fig 2.

Returning to our opening example involving racial profiling, the question we must ask is whether you or I have sufficient evidence upon which to believe that the black man standing outside the restaurant is a valet. We've already noted that the belief is morally problematic because it seems unfair. Nonetheless, it also seems like it's the belief you ought to have given the evidence. Notice that although Clifford states it as a moral imperative to believe in accordance with the evidence, moral considerations—such as whether the belief would be a racist belief—play no role in determining whether a belief is epistemically rational or not. Epistemic rationality is just a matter of the evidence you have and whether that evidence provides adequate justification for your beliefs. As Ben Shapiro might say, *the facts don't care about your feelings*. Translating this into our, and Clifford's, vocabulary, we might say that the evidence, or the reasons for believing, don't care about your feelings.

Now if that's the case: the question can be put like this. When we have these conflicts between fairness considerations and, on the other hand, what you should believe based on your evidence, what should you do?

Broadly speaking, the options you have available for answering this question range from either just accepting that this is an irresolvable dilemma to saying that there's no fact of the matter about what you should do and you should just choose fairness or accuracy. These options can be mapped as follows.

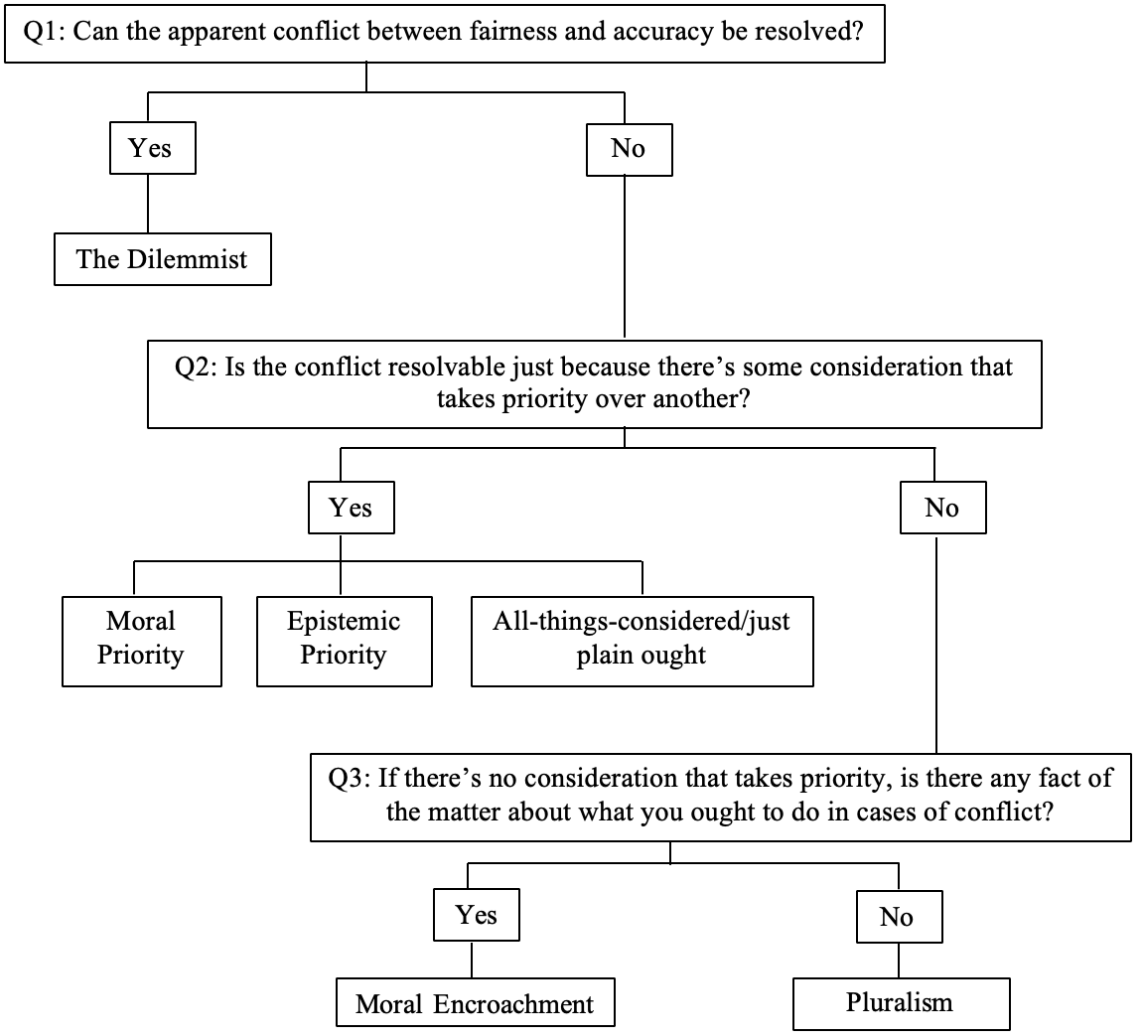

Fig 3.

Going forward, I'll begin by canvassing the various options—starting at the extremes with The Dilemmist vs. The Pluralist—in an attempt to present each camp *warts and all* so the reader can decide for themselves what they find most convincing. But I also think there is a right answer. My preferred answer, moral encroachment, requires a shift in how we understand the demands of fairness and accuracy. According to moral encroachment, there is no dilemma here, but there is a fact of the matter about what you ought to do. That sounds odd, but I'll try to show it's the most promising answer to this *apparent* conflict between fairness and accuracy.

2. The Dilemmist vs. the Pluralist

Dilemmas are familiar from our everyday lives. You want a decadent brownie but you also want to eat healthy. You can't do both. You promised a friend you'd meet them for lunch, but

then another friend has a crisis just before lunchtime. You either keep your promise to one friend or break the promise to go support your other friend. You love both your children equally but because Nazis are evil, they're threatening to kill both your children unless you choose one, and only one, for them to kill. These examples escalate quickly, but you get the point.

The possibility of genuine normative dilemmas is most clearly seen when both obligations are of the same normative kind, e.g., two moral demands. To illustrate this, let's consider Ginny. Ginny has promised Fred she'd see the new Marvel movie with him and only him (A), but she also promised George that she'd see the new DC movie with him and only him (B). Imagine that Ginny lives in a town with only one movie theater and because of her busy schedule she has been putting off her obligations until tonight, the last night that both movies are playing. Unfortunately, she can only either keep her promise to Fred and do A, or she can keep her promise to George and do B. There is no way for her to do both. Ginny finds herself torn between two incompatible options—A and B—and asks "What ought I do? A or B?" Let us now imagine that she shouts out this question to the universe and while she's shouting, a proponent of The Dilemmist position happens to be walking by. Noting Ginny's dilemma they helpfully answer, "Well, you ought to do A, and you ought to do B. Basically, you're just stuck between a rock and a hard place." We could reasonably expect Ginny to reply reminding this stranger offering advice that she can't do both, and she wants to know which she ought to do. If they were again to simply reply that she ought do A, and she ought do B, that is no help.

Now, why should we think that there are genuine dilemmas of this sort? Partly, we can follow a line presented by Bernard Williams (1965): in the case of moral dilemmas you might think that whatever Ginny does, she'll feel regret at not having done the other. If she breaks her promise to George and goes to the movies with Fred, she'll feel the need to make it up to George. Seeing that both incompatible acts are required does justice to why we would feel this regret. This case of promising involves two moral obligations, but what about when moral and epistemic obligations, such as fairness and accuracy, seem to collide?

A classic defender of this dilemmist position is Tamar Gendler. She introduces the following case that mirrors the conflicts we've been discussing so far, and argues that the characters in the case (taken from John Hope Franklin's autobiography) simply face an irresolvable dilemma.

> *Social Club.* Agnes and Esther are members of a swanky D.C. social club with a strict dress code of tuxedos for both male guests and staff members, and dresses for female guests and staff members. They are about to go on their evening walk so they head towards the coat check to collect their coats. As they approach the coat check, Agnes looks around for a staff member. All of the club's staff members are black, whereas only a small number of the club members are black. As Agnes looks around she notices a well-dressed black man standing off to the side and tells Esther, "There's a staff member. We can give our coat check ticket to him." (see Gendler 2011 for the original case)

Gendler, in a familiar refrain, argues that given that we live in a society structured by racial categories, we simply face a tragic irresolvable dilemma. We must either (a) lose out on knowledge and pay the epistemic cost of failing to attend to certain sorts of statistical information about cultural categories (i.e., encoding the information that a minuscule fraction of club members are black whereas all the staff members are black) or (b) we must actively believe against the evidence and work at regulating the inevitable implicit biases which that information gives us. In other words, we must choose between doing what we epistemically ought (attend and respond to the background statistical information about the race of the staff members) and what we morally ought (not use someone's race to make assumptions about them, such as that they are staff). This, she argues, places us in a tragic irresolvable dilemma. We cannot simultaneously fulfill both our moral and epistemic obligations, and there is no way to resolve this conflict. Both options have major downsides.

A challenge for interpreting these cases of normative conflict as genuine dilemmas is that it can undermine movements against racism and implicit bias. As Jennifer Saul (2018) argues, if we were to suggest that opposition to racism (what is morally required) leads one into (epistemic) irrationality, then the consequence is that the person committed to anti-racism will be irrational. As Saul (2018, pp. 238-9) goes on to note,

> Although it is clearly not Gendler's intent, this fits exceptionally well with the right-wing narratives of politically correct thought-police attempting to prevent people from facing up to difficult truths; and of the over-emotional left, which really needs to be corrected by the sound common sense of the right. Anything that props up these narratives runs the risk of working against the cause of social justice.

However, there may be ways to acknowledge that moral and epistemic obligations can conflict without suggesting that adopting the moral option means you're being irrational. One such way, pluralism, starts from the observation that obligations of all different sorts are in conflict all the time. Maybe, then, there's nothing special about moral-epistemic conflicts. According to *pluralism* there is a plurality of oughts and from the perspective of each ought you simply ought do what it prescribes. If we return to the flowchart from Fig. 3, the pluralist says that the conflict is resolvable, meaning that it's not the case that any one consideration (always, or even ever) takes priority over another, and there's just no fact of the matter about what you should do in these cases of conflict.

For example, consider the Knights Who Say Ni from *Monty Python and the Holy Grail*. To say that the Knights Who Say Ni have some weird norms would be an understatement. For those unfamiliar with this important cultural reference, one of the standards governing Knights Who Say Ni is to shout "Ni!" until their demands are met, i.e., a gift of shrubbery. As a result, that you are a Knight Who Says Ni generates a reason for you to shout "Ni!" until you are gifted shrubbery. Despite the silliness of the example, we recognize that a lot of prescriptive norms are like this. Perhaps the rules of etiquette only generate reasons for you —a reason to take off your hat when indoors, a reason to not pick your nose at the dinner table, etc.—if you care about the rules of etiquette. Perhaps these standards of fairness and accuracy are in the same boat.

A consequence of this style of pluralism is that if any standard can generate reasons for you, then we will have an infinite number of reasons generated by an infinite number of standards and we will constantly be pulled in every direction. Although this initially sounds bad, it does seem to reflect how we feel a lot of the time. Consider, for example, the familiar saying that you can either have it cheap, fast, or good: you can only pick two! Alternatively, you can either be well-rested, have good grades, or a social life, but not all three.

A challenge for the pluralist is that they need to explain why our intuitions in cases involving a small good in one domain at the expense of a large cost in another domain seem to suggest that there is something you *all things considered ought to do*. These cases are referred to as the *argument from notable-nominal comparisons* (see Chang 1997, Scanlon 1998, Parfit 2011). For example, if you are walking by a lake and see a drowning child you could either save the child or you could not. From the perspective of self-interest, perhaps you shouldn't save the child—after all, you would get your clothes wet and from the perspective of self-interest you prefer having dry clothes. From the perspective of morality, you should rescue the drowning

child. The pluralist in this case must simply say that what you morally ought do is rescue the child, but self-interestedly you ought not rescue the child. This seems like the wrong conclusion. What we want to say in this case is that you just ought to do what you morally ought do. Thus, we are pushed towards a different answer: the conflict is resolvable, so there must be some consideration that takes priority over the other considerations. So, let us now turn to those three options.

3. Moral Priority, Epistemic Priority, and the All Things Considered or Just Plain Ought

To get the intuitive grasp on the idea that there is an *all things considered ought* or *just plain ought* that can tell us what we should do when we're in a dilemma or we feel conflicted between competing options: consider the ought we deploy when we offer advice to one another, it is the kind of ought that would issue from a wise guru who has weighed all the relevant considerations. For example, in the movie adaptation of *The Notebook,* Noah (played by Ryan Gosling) asks Allie (played by Rachel McAdams) what she really wants. He doesn't want to know what it is that *everyone wants*, what *she thinks he wants* or what her *parents want*. Those wants are all *relative-wants*, they are wants relative to other people. Noah wants to know what Allie *just plain wants.* This example suggests that we can make sense of an idea of an all things considered ought, or just plain ought, and maybe determining what we ought *all things considered* do will help us in these cases of normative conflicts between fairness on the one hand and accuracy on the other.

Starting with the all things considered ought, one consideration in favor of such an account is that it offers us a unified and comprehensive answer. That is, the other oughts, such as the moral ought (as captured by considerations of fairness) and the epistemic ought (as captured by considerations of accuracy), are partial or incomplete collections of all the relevant considerations. The all things considered ought, as its name suggests, is comprehensive; it is based on *all* the considerations that weigh either in favor or against. We can contrast this with something like the moral priority or the epistemic priority view. According to those views, when Agnes (the swanky club member looking to drop off her coat) is deciding what to do or what to believe, either the moral considerations are more weighty and take priority over the epistemic considerations (moral priority, which means she should ignore the statistical fact that black people are almost entirely employees of the club, rather than members) or the opposite (epistemic priority, which means she should just go

with the statistics and ignore the potential harms for black club members like John Hope Franklin when others assume that they're employees). A challenge for going either of these two routes would be explaining why one takes priority over the other. Furthermore, this challenge is particularly difficult for anyone who wants to defend epistemic priority, i.e., that considerations of accuracy take priority over considerations of fairness.

To briefly expand on this challenge, it is simply not clear whether the epistemic ought is powerful enough to take priority over other considerations. According to an influential argument from Mark Nelson (2010), we have no *positive* epistemic duties to believe: that is, we are never *required* to believe anything. Evidence might give us a reason to believe something, but to say that it follows that I have a duty to believe everything for which I have evidence leads to a conclusion we ought reject: that I am required to believe an infinite number of things. For example, if you have evidence that supports believing p and you are required to believe p, then following simple rules of logic you now also have evidence that supports believe p or q and you're similarly required to believe p or q, you also now have evidence that supports believing p or q or r, etc. There are now infinitely many beliefs we are required to believe.

Returning to the all things considered view, we can instead envision some third perspective from which to adjudicate these competing demands and values. A benefit of the all things considered account is that it offers a kind of *comprehensive value* (see Chang 2003). Further, this more comprehensive value comes with a kind of special authority that the all things considered ought has that the other oughts and considerations don't. For every consideration, whether it is morality, epistemic rationality, etiquette, prudence, or whatever, we can always ask, "Why should I be moral?", "Why should I take my hat off when I enter a church?", "Why should I care about self-interest?", etc. The reason we can ask these questions of these oughts and considerations is because we are implicitly granting the authority of a more comprehensive ought, of a third perspective from which to answer these should questions.

However you wish to cash out this idea of an all things considered ought, there are some reasons for skepticism. Here I will canvass just two such reasons. First, there might be no common scale for weighing these considerations—moral, epistemic, aesthetic, legal, etc. That is, there is no further normative perspective from which we can both ask and answer the question of how these reasons should be combined. To see why one might doubt that there is a common scale for weighing moral and epistemic considerations together, consider

the following example: trying to compare different colleges to go to. There are many dimensions along which we can compare various colleges, e.g., average class size, professor to student ratio, sports teams, Greek life, or the best damn band in the land. But, how do we determine which college is *all things considered the best*? Similarly, imagine trying to give an all-things-considered judgment about who is the all-things-considered best superhero (it's Captain America btw, but I'm willing to hear arguments in favor of Thor). If we could answer that question, why do such arguments get so heated? If there were an answer, there wouldn't be reasonable disagreement. Maybe just as we can't make sense of the best *all things considered* college or superhero, we similarly can't make sense of the all things considered ought.

Second, and relatedly, you might worry that the all things considered or just plain ought doesn't make sense because it requires there to be some standard that is the most normatively important standard. However, either way you cash out this idea is incoherent. David Copp (1997) argues that whatever way we try to explain the authority that this most normatively important standard has will force us to embrace a contradiction. For example, consider the standard S. Let us suppose that S is the most normatively important standpoint. But, if S is the most normatively important standpoint, then there must be some *more* authoritative standard, A, that tells us that S is the most normatively important standpoint. The challenge we now face is answering whether A is identical to S. If A is identical to S, then A cannot play the role of the more important normative standard that establishes the normative supremacy of S. That kind of self-endorsement is characteristic of all the normative standards we've been considering and as such it is unimpressive, e.g., morality tells you to listen to morality and self-interest tells you to listen to self-interest. So, A cannot be identical to S. But, were A to be a standard *other than* S, that is similarly unimpressive. What we want to know is whether S is the most normatively important standard *full stop*. But if A is more authoritative, then S is not the most important standard. So, Copp concludes, there is no coherent way to cash out the idea that there is a normatively most important standard that is the standard of what we ought, all things considered, do.

To see this in play in Gendler's original example involving John Hope Franklin, suppose you think that morality is the normatively important standpoint. From that it follows that considerations of fairness take priority over considerations of accuracy. However, what makes morality the most normatively important standpoint? The answer can't be morality itself, because that's an unimpressive kind of self-endorsement characteristic of all normative standards. So, perhaps there's some *more authoritative* standard

than morality, and that more authoritative standard, The Authority, tells us that morality is the most normatively important standard. But, now it turns out that morality isn't the most normatively important standard, it's The Authority. Also, what makes The Authority the most normatively important standard—some other authority, The Meta-Authority?

4. Moral Encroachment

As I noted earlier, I don't personally find any of the previous options fully convincing. What leads those options astray, I believe, is the very framing of the scenarios under consideration as ones involving a *conflict* between *competing* norms. Rather, what I believe the valet case and the John Hope Franklin case present us with is just a morally fraught example of a very common occurrence: a high-stakes situation. There are many situations in which we would characterize what we should believe or what we should do as being *high stakes*. For example, consider the following classic example. Suppose you have an important mortgage payment due on Monday. As a millennial, the idea of a mortgage payment sounds like a luxurious expense I'll never have the opportunity to experience, so to update the luxurious expense example to one that I expect more readers will be able to relate to: imagine that you have a meal-kit delivery subscription that automatically withdraws a certain amount from your bank account every Monday.

Now, let's say that it's Friday afternoon, you've just been paid, but for some reason you don't have direct deposit set up. You have two options: you can either try to go to the bank on your way home on Friday or hope that the bank is open on Saturday and find some time to go then. Alternatively, you can deposit your check through the mobile-banking app on your phone. If you do the former, the money will be in your account right away (either on Friday or on Saturday), if you do the latter you know the check probably won't clear until Tuesday. Now imagine a low-stakes case: you've got more than enough money in your bank account so you don't risk overdraft fees if you don't deposit your paycheck before Monday. As you and your wife are on the bus headed home you're about the pass the bank and have to make a decision about whether to request the stop or stay on the bus to get home earlier. It's been a long day and since you're in no rush to deposit your paycheck, you causally ask your wife, "Hey, do you know if the bank is open on Saturday?" She tells you that she thinks she's been by the bank before on a Saturday and it was open.

Now contrast that case with a high-stakes case. You will face overdraft fees on Monday if the check is not deposited before Monday. If you decide not to stop in at the bank

on Friday and you decide to believe that your wife knows that bank will be open on Saturday and it turns out it's not open...bad news bears. Here's the thing, though: in both the low-stakes case and the high-stakes case, your evidential situation is the exact same: you are relying upon the testimony of your wife. If all that matters to epistemic rationality is believing on the basis of your evidence, well, you have the same evidence in both cases so what you should believe in both cases should be identical. Our intuitions, however, differ. In the low-stakes case it seems that your wife's testimony is sufficient evidence to be justified in believing that the bank will be open on Saturday, but in the high stakes case that very same evidence no longer seems sufficient. After all, if she's wrong then you risk overdraft fees.

These sorts of examples have been used to argue for *pragmatic encroachment*, the idea that practical features of our situation can make a difference to (or "encroach" upon) whether we're in a position to know or whether we're in a position to believe on the basis of the evidence we have available to us (Fantl and McGrath 2002 and 2009; Stanley 2005; and Schroeder 2012). However, recall Clifford's claim that "it is wrong always, everywhere, and for anyone, to believe anything upon insufficient evidence." What the pragmatic encroachers add to this observation from Clifford is that what counts as sufficient or insufficient evidence can vary according to the practical stakes of the belief.

If we return now to the cases of the mistaken valet and John Hope Franklin, I believe these cases are similar to the cases that motivate pragmatic encroachment (and I'm not alone, see also Code 1987, Pace 2011, Moss 2018a and 2018b, Basu and Schroeder 2019, Fritz 2018, Bolinger 2018, Basu Forthcoming). The moral risks present in the cases, i.e., the moral consideration that our belief that John Hope Franklin is a staff member would be unfair to John Hope Franklin for reasons previously discussed, makes the cases high-stakes cases. Moral encroachment, then, understood as the thesis that morally risky beliefs raise the threshold for (or "encroach" upon) justification makes sense of the intuitive thought that racist beliefs require more evidence and more justification than other beliefs.

What distinguishes this approach from others considered so far is that it provides an alternative route to simply throwing our hands up in the air and proclaiming these cases to be dilemmas. It also avoids the problems associated with moral and epistemic priority, because on this view neither consideration takes priority; rather, both considerations work together to determine what you should believe and what you should do. The moral considerations raise the epistemic standards in these high-stakes cases. We can preserve the thought that the facts don't care about your (or other people's) feelings while also

recognizing that whether or not you are justified in believing on the basis of the evidence available to you is a question that is sensitive to non-factual or non-evidential considerations. Whether you have enough evidence to believe varies according to the stakes.

Returning to the courtroom, we see this intuitive thought in play. Criminal cases are high-stakes cases; that is why the standard the evidence must meet is higher than in civil cases. In a criminal case you must prove beyond reasonable doubt that the defendant committed the crime; in a civil case the standard is considerably lower. Similarly, if you are going to believe that someone is a valet on the basis of their skin color, although that might give you a lot of evidence, it's not *enough* evidence to make the belief justified. Given the high moral stakes, you must look for more evidence. See also Siegel (this volume).

Despite my preference for moral encroachment, this view also faces some challenges. First, moral encroachment risks compounding and contributing to the unfairness that John Hope Franklin, and other folks who are constantly mistaken for staff or "the help," experience. Moral encroachment recommends that, when it comes to beliefs about black men and other non-dominantly situated groups, the epistemic situations are high stakes. As a result, when it comes to forming beliefs about dominantly situated folks, our epistemic situations will be more free and less burdensome because we won't constantly be walking on *epistemic eggshells*. This, however, also seems unfair. To answer this challenge, I think it's a mistake to say that moral considerations only make it *harder* to believe; sometimes moral considerations might also make it *easier* to believe. For example, if we live in a world in which women's testimony, and the testimony of victims of sexual assault more generally (regardless of the gender or sex of the victim), is routinely discounted, then perhaps we have a moral burden to lower our evidential standards for believing the victim (see Fricker 2007 for more on these forms of testimonial injustice). However, if we go this route, this opens us up to more challenges: how do we determine when moral considerations cause the threshold to go up and when moral considerations cause the threshold to go down? Also, we shouldn't downplay the worry that, if we allow considerations of morality and knowledge to encroach upon each other, our standards of justification might be easily manipulated.

Related to this epistemic eggshell worry is a worry about the demandingness of moral encroachment. If one will fail to be justified in virtue of failing to appreciate the burden and risks that they impose on another, then almost all beliefs about other people—especially any belief about a person on the basis of their race or another protected class—are going to be high stakes. Moral encroachment, it seems then, is demanding. It requires moral

agents to be fairly sophisticated in recognizing when they should occupy this kind of moral standpoint. However, as I've suggested elsewhere (Basu Forthcoming), the moral encroacher should just bite this bullet. Morality *is* demanding. It should not be surprising then that a moral constraint on our epistemic practices would be similarly demanding. In our everyday lives and our day-to-day beliefs, we may often fall short of the moral and epistemic ideal. The ideal, however, exists as a standard we ought to strive to meet nonetheless. Similarly, consider Clifford's response to the objector who says that they're a busy man and that they can't possibly be expected to take the time and effort to make sure they never ever believe on the basis of insufficient evidence. To such a character, Clifford simply offers the following rebuke, "then he should have no time to believe."

This list of objections is not exhaustive (see Gardiner 2018 for more). Nonetheless, I sincerely believe that moral encroachment offers the best analysis of the cases in which fairness and accuracy seem to conflict. To then finally answer the question contained in the title, fairness does not require inaccuracy. Nor does our desire for accurate beliefs require that we disregard considerations of fairness. *Community's* Pierce Hawthorne isn't a righteous jerk who at least is aiming for accurate beliefs when he accuses Troy of stealing Annie's pen because Troy (a young black man) is the statistically likeliest candidate. Pierce Hawthorne isn't exhibiting any epistemic virtues. He's just being a jerk.

References

Basu, R. 2018. The Wrongs of Racist Beliefs. Philosophical Studies:1-19.

Basu, R. 2019. What We Epistemically Owe To Each Other. Philosophical Studies 176(4): 915-931.

Basu, R. Forthcoming. Radical Moral Encroachment: The Moral Stakes of Racist Beliefs. *Philosophical Issues.*

Basu, R. and Schroeder, M. 2019. Doxastic Wronging. In Pragmatic Encroachment in Epistemology, eds. Brian Kim and Matthew McGrath. Routledge. Pp. 181-205.

Beeghly, E. 2015. What is a Stereotype? What is Stereotyping? Hypatia 30(3):675-691

Beeghly, E. 2018. Failing to Treat Persons as Individuals. Ergo 5(6):687-711.

Blum, L. 2004. Stereotypes and Stereotyping: A Moral Analysis. Philosophical Papers, 33(3): 251–289.

Bolinger, R. 2018. The rational impermissibility of accepting (some) racial generalizations. Synthese.

Buchak, L. 2014. Belief, credence and norms. Philosophical Studies 169:285-311.

Chang, R. 2003. All Things Considered. Philosophical Perspectives. 18(1):1-22.

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

Clifford, W. K. 1877. The Ethics of Belief. Available online: http://people.brandeis.edu/~teuber/Clifford_ethics.pdf

Code, L. 1987. Epistemic Responsibility. Brown University Press.

Copp, D. 1997. The Ring of Gyges: Overridingness and the Unity of Reason. Social Philosophy and Policy 14(1):86-101.

Enoch, D. and Spectre, L., and Fisher, T. Statistical Evidence, Sensitivity, and the Legal Value of Knowledge. Philosophy and Public Affairs 40(3).

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

Fantl, J. and McGrath, M. 2009. Knowledge in an Uncertain World. Oxford University Press.

Fricker, M. 2007. Testimonial Injustice: The Ethics of Knowing. Oxford University Press.

Fritz, J. 2018. From Pragmatic Encroachment to Moral Encroachment. Pacific Philosophical Quarterly.

Gardiner, G. 2018. Evidentialism and Moral Encroachment. In Believing in Accordance with the Evidence: New Essays on Evidentialism, ed. Kevin McCain. Springer.

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

Kelly, T. 2016. Evidence. The Stanford Encyclopedia of Philosophy. Link: http://plato.stanford.edu/archives/fall2014/entries/evidence/

Lippert-Rasmussen, K. 2011. 'We are all Different': Statistical Discrimination and the Right to Be Treated as an Individual. Journal of Ethics 15(1):47–59.

Moss, S. 2018a. Probabilistic Knowledge. Oxford University Press.

Moss, S. 2018b. Moral Encroachment. Proceedings of the Aristotelian Society.

Munton, J. 2017. Perceptual Skill and Social Structure. Philosophy and Phenomenological Research.

Nelson, M. 2010. We Have No Positive Epistemic Duties. Mind 119.

Pace, M. 2011. The Epistemic Value of Moral Considerations: Justification, Moral Encroachment, and James' 'Will to Believe'. Nous 45(2):239-268.

Parfit, D. 2011. On What Matters. Oxford University Press.

Saul, J., 2018. (How) Should We Tell Implicit Bias Stories? Disputatio 10, 217–244. https://doi.org/10.2478/disp-2018-0014

Scanlon, T. 1998. What We Owe To Each Other. Harvard University Press.

Schauer, F. 2006. Profiles, Probabilities, and Stereotypes. Harvard University Press.

Schroeder, M. 2012. Stakes, Witholding, and Pragmatic Encroachment on Knowledge. Philosophical Studies 160(2):265-285.

Siegel. S. This Volume.

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

Williams, B. 1965. Ethical Consistency. Proceedings of the Aristotelian Society 39.