

Epistemic Justice and the Principle of Total Evidence

S. Roush

I have a dream that my four little children will one day live in a nation where they will not be judged by the color of their skin but by the content of their character.

– Martin Luther King, Jr., August 28, 1963, March on Washington for Jobs and Freedom

Abstract

Epistemic injustice is injustice to a person *qua* knower. In one form of this phenomenon a speaker's testimony is denied credence in a way that wrongs them. I argue that the received definition of this testimonial injustice relies too heavily on epistemic criteria that cannot explain why the moral concept of injustice should be invoked. I give an account of the nature of the wrong of epistemic injustice that has it depend not on the accuracy of judgments that are used or made in the process of deciding whether to listen to or trust a speaker, but on whether the basis of the decision about a speaker is their reliability or their identity, and the account explains why the latter is a moral wrong. A key difference between the two accounts is how they classify the use of true statistical generalizations connecting identity and reliability. The received view implies that this cannot be an injustice, while the view proposed here implies that it can. As such the new view appears to imply a conflict between moral and epistemic obligations: it is morally wrong to use true statistical generalizations in certain contexts, yet they are part of our evidence, and we are epistemically obligated to take all of our evidence into account. I reconcile these two thoughts without adopting the currently popular view that a belief's being morally wrong makes it epistemically unjustified, and I argue that following the principle of total evidence encourages epistemic justice rather than thwarting it.

I. Introduction

II. What is the wrong of Epistemic Injustice?

III. Employment Justice: What does a job applicant deserve?

IV. Epistemic Justice: What does a speaker deserve?

V. Frequency Statistics and Epistemic Obligations

VI. Delayed Decisions – Waiting for Evidence

VII. Conclusion

I. Introduction

Epistemic injustice is injustice toward a person *qua* provider of knowledge. The phenomenon is often easy to recognize intuitively, and has wide-ranging ramifications. (Fricker, M. 2007, Kidd et al. 2017) If we treat women's testimony as worth only half as much as men's as a matter of principle and policy, for example, then many women will not be accorded the credibility they deserve, and the interests of women and of getting at the truth will suffer. Epistemic injustice is rarely as explicit or deliberate as in this example, but its unconscious and ingenuous forms are commonplace; new examples are now described regularly. Understudied are the questions what we should be aiming for if we aim to achieve epistemic justice, and what is the nature of the wrong we want to avoid.¹ Identifying the positive goal can help us to better understand its lack and can have an ameliorative effect, and characterizing the wrong differently can yield a different extension of the concept.

Miranda Fricker identifies the wrong of testimonial injustice both etiologically and counterfactually. On her account it is caused by use of unreliable stereotypes, and it involves the hearer giving the speaker "less credibility than he would otherwise have given". (Fricker, M., 42) This direction of definition is natural because of the psychology of clear cases of epistemic injustice and because in the motivating cases a speaker has clearly gotten less attention than they deserved. The restriction of the scope of injustice to the use of stereotypes that are unreliable (i.e., not mostly true) is a way to prevent too many things from counting as injustices. It is also motivated, I will argue, by an assumption of harmony between our epistemic and moral obligations, and supported by a failure to fully consider the fact that negative correlations between identity and abilities do exist. I will argue that the definition of epistemic injustice resting so heavily on the unreliability of the stereotypes that so often lead to it also makes out the phenomenon to be essentially an epistemic mistake rather than a moral wrong, leaving it unclear why the concept of injustice is relevant.

Below I propose an alternative account of epistemic injustice that is based on an analogy with US employment discrimination law. In this legal framework making negative employment decisions on the basis of protected identity traits is forbidden unconditionally, regardless of whether there are population-level correlations between identity traits and job qualifications. I define notions of fairness and justice in decision-making that explain why the law should take this form, and develop an analogous set of requirements for testimonial justice. On these views, injustices can occur even if the identity-generalization a decision is based upon is true, and even if the person denied a job or attention also would not have gotten it had they been judged by their abilities. The wrong resides in the basis of the decision, not in its outcome or the accuracy of beliefs that led to it.

¹ Fricker (2007, Chapter 4) gives an account of testimonial justice as a virtue of the hearer. The question here is about the state or benefit we would be achieving for the speaker, but I will make comparisons between the two views below.

Prohibiting the use of true statistical generalizations seems to run afoul of requirements that epistemologists of many persuasions hold dear, principles that exhort us to take all of our evidence into account, not just some of it. How can ignoring evidence be epistemically permissible? There are two distinct principles that have run under the banner of “total evidence”. One, usually referred to as the Requirement of Total Evidence, demands that in our beliefs we take into account all of the evidence that we currently have. The other, usually referred to as the Principle of Total Evidence, says, with important conditions, that we will be better off acquiring and assimilating more new evidence than we would be if we didn’t. The first principle, I argue, is difficult to justify, and the only cogent justification of an unconditional principle of this sort doesn’t apply to naked statistical evidence (frequency data).

Naked statistical evidence can justify belief states about individuals, but no expert thinks this comes without some further condition(s). Though the strongest version of such a condition defeats many everyday profiling inferences, it is possible to meet even that condition in the context of employment decisions, and conceivably in speaker-hearer contexts. Thus, the most careful epistemic considerations about statistical inference do not uniformly prohibit us from having profiling belief states – such as a very high confidence that candidate #29 is not an orthopedic surgeon on the basis merely that so few women are and #29 is a woman. Such belief states seem to support actions that I argue here are illegal or immoral, or both, suggesting a tension between epistemic and moral considerations. I argue that this need not be seen as a conflict, because belief states and actions are two different things, and the former do not automatically lead to or justify the latter.

In the expected utility paradigm belief states don’t justify any action unless they are combined with a subject’s utilities, or values. Moreover, there is a further choice point between beliefs and utilities on the one hand and undertaking the action that they justify on the other. The second total-evidence principle – the Principle of Total Evidence – concerns this choice, which is whether to act on our current beliefs or postpone action while we gather more evidence and potentially come to different beliefs on which to base our choice of action. In contexts where justice is paramount, this principle supports not acting (e.g., deciding whether to hire or to trust) until we gather more specific evidence than a statistical generalization can give. This analysis shows that it is possible to meet the moral requirements of epistemic justice without violating epistemic requirements, and also without conflating the two. To have a high credence that #29 is not an orthopedic surgeon is not the same thing as rejecting her job application. And to refrain from rejecting #29’s job application while you read her file is not to ignore statistical evidence. What disqualifying #29 from the competition without looking in her file indicates is not stalwart adherence to the evidence but that the decision-maker considers the cost of reading the file to be greater than the cost of an unjust decision.

II. What is the Wrong of Epistemic Injustice?

Epistemic injustice can take many forms, but in the case of testimonial injustice that is my focus here, Fricker describes what the cases have in common this way: “prejudice on the hearer’s part causes him to give the speaker less credibility than he would otherwise have given.” (Fricker, M. 2007, 42) In testimonial injustice, prejudices – pre-judgments that are resistant to counter-evidence owing to some affective motivation, that may or may not be ethically bad – combine with stereotypes – widely held generalizations associating certain social groups with certain traits – to produce a credibility deficit for the speaker in the eyes of the hearer. For injustice, the stereotype must be unreliable and its use must be disparaging, and an identity-prejudicial stereotype is negative when its resistance to evidence comes from motivation that is ethically bad. Identity prejudices and stereotypes that lead to injustice distort the hearer’s perception of a speaker’s credibility, without necessarily involving a mental state of belief in the associated generalization. That is, even when one’s beliefs are enlightened, one’s social perceptive reflexes may not have caught up. (Fricker, M. 2007, 32-39)

On Fricker’s use of the concept of stereotype, the generalization that a stereotype embodies may be reliable or unreliable. On her view reliable stereotypes are an essential heuristic for making credibility judgments in everyday testimonial interactions. In order to be a negative identity-prejudicial stereotype, and so in order to lead to injustice, on her view, the generalization expressed in a stereotype must be unreliable – i.e., not mostly true – and (the prejudicial aspect) maintained without proper regard for evidence.

I will reject the claim that a generalization must be unreliable in order for its use to count as an injustice, but I recognize several possible sources of pressure toward this view. One is a concern about over-inclusiveness; if stereotypes are essential heuristics for testimonial credibility judgments we make on a daily basis, as Fricker thinks, then if even those based on reliable generalizations caused injustice, there would be injustice always and everywhere, potentially trivializing the notion. I will return to this concern below.

Even a reliable generalization may have exceptions, so its use may lead to a false belief about a speaker’s credibility. But, according to Fricker, because we know that most used car salesmen are untrustworthy, I do not commit an injustice by not trusting the one before me about whom I have no specific evidence, even if he is the (supposedly) exceptional honest man of the group. (Fricker, M. 2007, 42) This is a natural response to a second source of pressure to deny that use of reliable generalizations is an injustice: the idea that it is unfair to a subject to take her as epistemically culpable when she has followed the available evidence in forming her belief. This in turn is supported by the sound epistemic fallibilist thought that we can’t require a person to do so much that she is guaranteed to have gotten a true belief, in order to be counted as justified in having the belief. Justified belief is compatible with the possibility of bad luck that makes the belief false; otherwise we would have no justified empirical beliefs about the world at all.

On epistemic grounds, bad luck seems to excuse the hearer of epistemic blame when she uses a reliable, well-evidenced generalization about members of a social group to come to a false

assessment of one of its members who is exceptional. But this is then apparently taken by Fricker to be sufficient to show that the (supposedly) rare honest used-car salesman has not been done an *injustice*:

... [T]he hearer has not put a foot wrong—she has made a credibility judgement that is in line with the evidence, ... [It is a case] of innocent error on the part of the hearer: no epistemic culpability, and no ethical culpability. There is no testimonial injustice here, and our definition stands.” (Fricker, M. 2007, 42-43)

The honest used-car salesman has not been accorded the credibility he would have gotten had he not been judged by means of a stereotype, in accord with part of Fricker’s definition of injustice cited above, yet it does not count as an injustice to him but as mere bad luck for him just as it is epistemic bad luck for the hearer.

This is a hearer-centric assessment in which we fail to consider the situation from the honest used-car salesman’s point of view. For an indication in favor of the view that I will ultimately take on the issue of well-evidenced, reliable generalizations, notice that the bad luck of the hearer and the speaker in such a case are not symmetrical. The hearer has made an epistemic error that may deprive him of information, but the speaker has taken a loss that seems unfair; he is not trusted as a witness, even though he is trustworthy, not because of any trait or action individual to himself, but because he can be classified as a member of a particular group. If one is not moved by the plight of honest but distrusted used-car salesmen, that may be because *used-car salesman* isn’t the right kind of identity trait for justice to be at issue rather than because use of reliable generalizations is always innocent. I will discuss more rousing examples below.

A third pressure to take Fricker’s position on reliable generalizations comes from an assumption that there is harmony between our epistemic and moral obligations and results. This in turn may be motivated by the thought that ought implies can, and that we cannot fulfill conflicting obligations. The expectation that being epistemically rational in this domain will automatically fulfill our moral obligations is evident in Fricker’s (2007) discussion: “Ultimately, the point is to see how our epistemic conduct can become at once more rational and more just.” (Fricker 2007, 4) “Epistemological nuance aside, the hearer’s obligation is obvious: she must match the level of credibility she attributes to her interlocutor to the evidence that he is offering the truth.” (Fricker, M. 2007, 19) I don’t take it as obvious a priori that if our decisions about which speakers to trust fulfill our epistemic obligations then they ipso facto fulfill our moral ones.² I will investigate our ethical and epistemic obligations here separately, as far as possible, to discover how they may relate. I will allow for the possibility that our situation is, for want of a better term, tragic, where it is sometimes not even possible to fulfill both our epistemic and our

² There are arguments for aspects of such a view, e.g., Rima Basu’s (2018) argument that no one could be a perfectly rational racist, but I will not assume their conclusions here.

ethical obligations at the same time.³ More mundanely, the situation might present a trade-off between ethical and epistemic benefits.⁴

In defining epistemic injustice Fricker invokes epistemic considerations in three distinct ways. One is in identifying the aspect of the speaker that is the object or site of the injustice; it is her capacity as a knower. A speaker may suffer harms of being insulted, undermined, or humiliated, but these are secondary. Since her capacity as a knower is central to her humanity, the injustice she suffers from disrespect of that is intrinsic. (Fricker, M. 2007, 44) A second role for epistemic considerations is in the fact that a hearer's act of coming to regard someone as a knower or not, is done, whether consciously or unconsciously, whether to form a belief or to form a perception, on the basis of some information, and it is done in a way that promotes an accurate picture or does not, that is, in an epistemically sound or unsound way.

One can imagine that epistemic facts about the hearer's formation of attitudes have a role as a means to just and unjust results, without the epistemic facts being defining or determinative of whether justice or injustice results. But the third way that epistemic considerations enter Fricker's account is in determining or defining the unjustness of a hearer's act. Fricker definitely intends for this injustice to be a moral wrong, and it needs to be in order to warrant the term "injustice" rather than "mistake", but the definitions she gives are at risk of making the wrong epiphenomenal on epistemic failures, firstly because the epistemic failures appear to be the only wrongs that are present in all cases of this injustice.

To see this, consider the candidates for explaining the moral wrong of epistemic injustice. The formulation that says epistemic injustice is disrespecting a speaker in her capacity as a knower appears as part of the effort to say what the wrong of epistemic injustice is, as distinct from how it is generated psychologically. This locates the object of the injustice, and what is apparently supposed to make the offense an injustice is that what is disrespected – her capacity as a knower – is part of her very humanity, rationality being distinctive of humanity. This still leaves it somewhat unclear what the nature, as opposed to the target, of the moral wrong is, but there is a more pedestrian problem. There is a great deal of testimonial injustice in which one's capacity as a knower full stop is not in question, in which what is dismissed is not the speaker's capacity as a knower of X for any X, but only as a knower for some X.

I might dismiss the possibility that Elaine could be good at math because she is female. That would be an injustice, but it would not be a denial of the capacity of knowing that has a part in making her human.⁵ I might believe that Elaine knows lots of things. I might even think it is likely that she is good at learning languages, because of the complementary stereotype about

³ For a view that makes it possible to satisfy epistemic and ethical obligations simultaneously, see Moss (2018, ch. 10). However, this comes at the price of embracing moral encroachment, the view that the epistemic status of a belief depends in part on moral factors, a view that I will not assume or endorse. See below, section V.

⁴ Authors who see a trade-off in this area are Kelly and Roedder (2008, 350) and Gendler (2011, 57).

⁵ If being good at math were part of our very humanity, then many we would have thought were human would have to be re-classified.

being female. A testimonial injustice may not rise to the level of being a denial of part of the speaker's very humanity. It may simply unjustly dismiss her authority on X. Disrespect of one's capacity as a knower full stop is not necessary for testimonial injustice.

Fricker concedes that Tom Robinson, the Black man in *To Kill a Mockingbird* (Lee 2010) who is on trial for the rape of a young white woman, and whose testimony is not respected, would nevertheless be trusted by the same people in knowledge exchange on other topics, those having to do with his work, for example. To see what makes the jury's rejection unjust in that case, she says, we have to look at the power structure and dehumanizing oppression of the social context in which the events of his trial take place. (Fricker, M. 2007, 130-131) But not all selective dismissals of a person's expertise or truthfulness are committed against people below the hearer in an oppressive hierarchy. A student might sneeringly dismiss the possibility that a teacher knows any mathematics on the basis of the teacher being a woman. While dismissal of someone's capacity as a knower full stop and participating willingly in social oppression are plausibly moral wrongs, neither is essential for epistemic injustice.

Another form that disrespect can take is objectification, treating the speaker as a possible source of information but not as a participant in the activity of pooling knowledge that human beings do with each other. (Fricker, M. 2007, 132) I might figure out what time it is by noticing that a particular person went to lunch, but not regard them as someone whom I would deign to ask any question. This definitely seems to qualify as a moral kind of wrong, but again it does not seem to be necessary to objectify people in order to commit an epistemic injustice. I don't objectify Elaine when I presume she is bad at math and talented at languages; I regard her as someone who could teach me another language, and not merely as an object that indicates information.

Another possible source of moral explanation of what makes the epistemic kind of disrespect an injustice is a property of the motivation behind the choice to mistrust a speaker or to ignore evidence about them, namely that the motivation was malicious or otherwise ethically bad. It is part of Fricker's definition of a negative prejudicial stereotype that its resistance to counterevidence has an ethically bad affective investment, but it is hard to adhere to this rigidly in cases where the perpetrator is trapped in an oppressive time period: she thinks that in *The Talented Mr. Ripley* (Minghella 1999), the active exercise of identity power in Mr. Greenleaf's dismissal of Marge Sherwood's reasoning as female intuition might be "well-intentioned and benevolently paternal" (Fricker, M. 2007, 14-15); it is nevertheless an injustice. I am also not convinced that bad intentions or affective investment are automatically sufficient to make something a moral wrong, especially an injustice, or that bad intentions could be exhaustive of what makes the wrong of epistemic injustice a moral matter.

For understanding of the moral wrong we might look to what is required for a stereotype to be negatively identity-prejudicial. For one thing it must be disparaging, but merely attributing a negative property to someone is not necessarily a moral wrong. Even if the person doesn't have the property, the ascription may simply be an epistemic error. Another facet of what happens

in the epistemic kind of injustice is that the victim is, in some sense, accorded less credibility than is her due, but again this could be due to an epistemic error on the hearer's part. Also, though having less credibility than one is due is usually a negative consequence, a loss and often a harm, harming is not automatically sufficient for wronging. Perhaps the moral wrong comes from how we come to the credibility assessment; when we use a negative prejudicial stereotype we act toward the speaker *qua* member of a social group, rather than as his own person. (Fricker, M. 2007, 35) But we judge the honest used-car salesman negatively using a stereotype belonging to a social group, and Fricker doesn't take this to lend any moral wrongfulness to the unlucky epistemic error, because, supposedly, the generalization we used was reliable.

The only thing left, it seems, to determine that use of a disparaging stereotype leads to a moral wrong of injustice, and that will be present in every case of the injustice, is what is supposed to make it prejudicial: its epistemic culpability, identified as unreliability of the generalization and resistance to counter-evidence.⁶ (Fricker, M. 2007, 32) As the only features that turn out to be necessary for the injustice, these properties are jointly sufficient, and thus defining of the concept. These are epistemic properties, so it is hard to see the offense as other than an epiphenomenon of an epistemic error, one that may or may not be accompanied by any of an assortment of moral wrongs. There is a formal solution to this problem, which is to make epistemic injustice a moral wrong by requiring that one or other of the moral wrongs canvassed above be present in a case for it to count as an injustice. However, this is merely formal unless we can explain why the wrongs just discussed should be called the particular kind of moral wrongs that count as injustice.

In what follows I give an account of epistemic injustice that explains why it is an injustice, a moral wrong. It is inspired by an analogy with US law on employment discrimination, in which I take attending to, heeding, or trusting a speaker to be like granting a speaker the job of providing you with knowledge. The account is based on a concept of fairness, that people should be judged in the same way, with the same results, as those with equal qualifications (or reliability). This yields an account of just decision-making in terms of the *basis* of decisions, namely that candidates (speakers) be judged on the basis of qualifications (reliability). This picture can explain why it seems wrong to judge a person *qua* member of a social group. It also explains why it seems wrong to be accorded less credibility than one would have been given but for the use of a naked generalization, although we will also see that the latter is too strong a requirement for a decision to count as unjust. In this account there is no need to appeal to epistemic considerations in describing the nature of the wrong. In particular, we don't need to

⁶ Resistance to counterevidence may be morally culpable, a possibility I allow below, but even so it would not explain why the wrong is *injustice*.

appeal to the unreliability of the generalization. We have moral reason to avoid using naked generalizations about some matters, whether the generalizations are reliable or unreliable.⁷

III. Employment Justice: What does a job applicant deserve?

The view that we don't commit injustice if the generalizations about identity that we use to judge peoples' qualifications are well-evidenced and reliable stands in contrast to the way American employment law handles discrimination. Title VII of the US Civil Rights Act of 1964 says, among other things, that

It shall be an unlawful employment practice for an employer

(1) to fail or refuse to hire or to discharge any individual, or otherwise to discriminate against any individual with respect to his compensation, terms, conditions, or privileges of employment, because of such individual's race, color, religion, sex, or national origin; ... (Title VII 1964, §703)⁸

This law is unconditional in that it prohibits negative decisions based on any of the protected identity characteristics regardless of whether any generalizations used are reliable or well-evidenced. The generalization that men are on average physically stronger than women is true and well-known, but hiring a candidate on the basis of this generalization and his being a man would be unlawful. If physical strength is relevant to a job, then one is expected to evaluate the individual's physical strength. Failure to hire a woman because she is a woman and the job requires physical strength is epistemically analogous to not trusting a used-car salesman because he is a used car salesman and they are (supposedly) generally dishonest. The difference in our reactions to the two cases, if there is one, tracks something of moral significance about legally protected identities.

One might hope that correlations such as that between sex and physical strength are rare, that use of identity would at least for the most part be an epistemic mistake of using an unreliable generalization, making the distinction between Fricker's approach and discrimination law something of a quibble. But there is an abundance of well-evidenced generalizations⁹

⁷ That the moral wrong of epistemic injustice is not itself an epistemic error, and that it can be committed without making the epistemic errors that Fricker identifies, do not of course preclude the possibilities that there are many epistemic errors people commit using generalizations – there are – and many epistemic deficiencies one might be able to identify in true naked statistical generalizations – there are. (For the former see Gardiner 2018; for the latter see, e.g., Thomson 1986, Wright 2011, Enoch and Fisher 2015, Pritchard 2018, Smith 2018.)

⁸ Discrimination on the basis of age and disability is prohibited by other statutes: the Age Discrimination in Employment Act of 1967 and the Americans with Disabilities Act of 1990. These statutes use language similar to Title VII, all referring to the basis of a decision.

⁹ I am taking well-evidenced frequency statistics (mere correlations between traits) to be well-evidenced reliable generalizations. There are many kinds of generalizations – e.g., generics, and relations between averages of populations – but frequency statistics are the kind of generalizations that concern me because they are so readily available.

connecting expertise with identity characteristics that have at least the appearance of relevance to judgments of a job candidate's qualifications. As of 2017, in the US, of orthopedic surgeons who reported their sex, 5% were women, 95% men¹⁰ (AAMC 2018), and this is not an isolated type of case. One doesn't need to be an official at the Bureau of Labor Statistics to be well-acquainted with the broad outlines of many of the realms of such correlations on which they collect data. No American would be too surprised to learn that 85% of lawyers in the US self-identify as white¹¹ (Cassens Weiss 2018), or that 95% of civilian aircraft pilots and flight engineers in the US are white (US Bureau of Labor Statistics 2018, 64). A fan of American football wouldn't be too far off if she guessed what percentage of NFL quarterbacks are white (78.8%, Gertz 2016), or what percentage of NFL players are white (26.5%, Gertz 2016).

Counting the use of reliable identity-based generalizations to judge a speaker's reliability as an injustice might threaten to turn many of our everyday testimonial interactions into ethical failures. But not counting the use of reliable identity-based generalizations as an injustice in employment decisions would contribute to perpetuating unjust inequalities. To see the direction of contribution that use of current correlations would always have, imagine we had used only the de facto correlation between white race and quarterbacking in 1970 to make hiring decisions at that position. If we took that correlation as our only evidence of the level of talent Black men have for the position,¹² then the percentage of Black quarterbacks would have remained vanishingly low. Even if combined with other evidence, use of the correlation would make a contribution in the direction of maintaining the status quo distributions of identities over jobs.

To understand why use of reliable generalizations can cause injustice, consider the following employment decision scenario. A hospital has two job openings for doctors, one in orthopedic surgery, the other in pediatrics. Suppose the percentage of orthopedic surgeons who are women is 4%, as it was in 2010 (AAMC 2012)¹³, and suppose, contrary to fact, that the proportions of the sexes among pediatricians are exactly opposite that among orthopedic surgeons – that only 4% of pediatricians are men.¹⁴ Suppose that the hospital has received 50 applications for each of these two jobs, and that the separate computer folders where they had been saved got accidentally but irreversibly merged.

¹⁰ In 2017, of the 861,770 active physicians in the census, who were offered the options of male, female, and no answer, 1,104 did not answer. See Rohde 2016 for discussion of gender in orthopedic surgery. Among residents in 2018-19, 15.4% identified themselves as female. (Brotherton and Etzel 2019)

¹¹ This is an estimate based on 20 states reporting.

¹² Of course in the real history the stereotypes of Black men as lacking leadership abilities were independently sourced, causative of and not merely caused by the low representation of Black men in this position.

¹³ That is, among those orthopedic surgeons who reported their sex, 4% reported female, 96% reported male. Out of 798,398 total physicians reporting, who were offered a choice of male, female or no answer, 1,086 did not answer. (AAMC 2012) I regret the gender-binary quality of this example, but male/female/not-indicated are the only data that were collected, and the specialties of those physicians not indicating their sex are not reported.

¹⁴ The actual percentage of women among pediatricians who reported their sex as male or female in the US in 2017 was about 63%. (AAMC 2018)

It would be laborious to go into every applicant file to determine which job they had applied for, but the files have been electronically labelled by sex (for subsequent compiling of equal employment opportunity statistics). Your day as an administrator would go much faster if you applied a sex filter to the merged folder, putting all the files labeled male into a new folder for the orthopedic surgery position, and all the files labeled female into a folder for the pediatrics position, and relaying these folders to the respective departments of the hospital. This sorting method would be highly reliable – we can expect that a mere four people out of the 100 would be misclassified – but, assuming there are no pediatric orthopedic surgeons in the applicant pool, and assuming there’s no correction down the line in the process,¹⁵ these four people would not be hired. I take it intuitively that these four people will have been treated unjustly, via the act of filtering by sex, an intuition that matches Title VII law.

Imagining the gut reaction those four people would have on learning what had been done is an initial guide to what’s gone wrong: that was *not fair*. For all you knew they were at least as qualified as the other candidates for the job they applied for, but they were excluded from consideration on the basis of something that wasn’t a qualification; they weren’t even given a chance to compete. To see that this is the important detail, imagine that those four candidates were actually not as well-qualified as the others who applied for the job they applied for. This would make their not getting the jobs a fitting outcome, but it would do nothing to alleviate the unfairness in the way they’d been treated. Title VII is in agreement.¹⁶ That the distribution of job results and job qualifications happen to match in a given case is not sufficient for equality of the treatment.

A central part of what is unfair about the filtering procedure is that it allows equally qualified candidates to be treated differently in the evaluation. All four of the applicants might have been less qualified than the rest of the pool for the job they applied for, but any of them also might have had qualifications equal to or higher than some who made it through the sex filter. The filter guarantees that two candidates for the orthopedic surgery position who were equally qualified but of different sexes would be treated differently, having different chances of getting the job. Likewise for the pediatrics position. The filter discriminates, intuitively and according to the law, even though you didn’t have any ill intent when applying it – you did it to save time and effort.

¹⁵ Suppose that the committees don’t know of each other’s searches, and committee members are very busy, so on seeing, for example, a pediatrician’s file among the applications for the orthopedic surgery position, they infer that the applicant has made a mistake and ignore the file rather than trying to re-direct it.

¹⁶ Supposing those four weren’t qualified, the employer would be able to show that she had a business interest in not hiring them, and that they wouldn’t have gotten the jobs even if the filter had not been used (a “same-decision” defense). The filtering would not have been a “but-for” cause, and the applicants would not be awarded jobs or back pay as damages. However because the filter made sex part of the basis of those employment decisions, and was used knowingly, the employer would still be liable for discrimination, and subject to injunctions and payment for things such as the applicants’ attorney fees and court costs. (Sullivan 2009, 199; Friedman 2017, 15-16) This point about liability in mixed-motive cases was entered into statute in the updated Civil Rights Law of 1991, §703 (m).

We can express a fairness condition that would prevent this scenario by saying that people with equal qualifications and in similar circumstances¹⁷ should be treated the same as to whether they get the job:

For all applicants x and y , for q a measure of properties needed for job J , and s a measure of circumstances,

$$q(x) = q(y) \text{ and } s(x) = s(y) \rightarrow (J(x) \text{ and } J(y)) \text{ or } (\text{not-}J(x) \text{ and } \text{not-}J(y)) \text{ or } P(J(x)) = P(J(y))$$

Fairness

If x and y had the same relevant qualifications and were in similar circumstances – e.g., same job in question, similar working conditions – then either both x and y would get a job or neither would get a job, or the two would have equal probabilities of getting a job. A person would be treated the same as any of her qualification-twins, with respect to getting hired or not.

This condition and the rest of what I have to say about fairness in employment are not citations or interpretations of Title VII's requirements. But my account is inspired by this law, and my conditions on decision-making justice, could, I think, make sense as part of a justification of this law, so I will make comparisons where appropriate. *Fairness* goes beyond Title VII by naming what the basis of hiring should be – qualifications – rather than only what it should not be – protected identity characteristics.¹⁸ This addition is helpful for sketching a positive conception of employment justice, some advantages of which will be seen below. Notice that, as in our

¹⁷ Circumstances include, for example, the job in question. One role of circumstances in the conditions I am defining is to enable us to make distinctions among different cases. A person may discriminate on the basis of sex for one type of job but not for another. Another role is to prevent fairness from requiring too much, namely that the employer would do the right thing in all possible circumstances. It is the actual circumstances that matter to the evaluation of a given case. Another is that circumstances can make a hiring decision legal that would not in other circumstances be, as when there is a business interest in the employee having a trait.

¹⁸ In taking qualifications to be the only legitimate grounds for decisions I am not only going beyond but also simplifying in a way not consistent with Title VII, for there are legal, non-discriminatory grounds for employment decisions that are not qualifications. It is legal, for example, to fail to hire an able-bodied person because he says he will not play intramural basketball in the after-hours activity the company uses for morale-building. (Friedman 2017, 15-20) However, if use of a trait that is not a protected identity trait leads to outcomes correlated with protected identity traits (disparate impact), then the trait must be related to job performance or supported by business necessity. (Friedman 2017, 41ff.) For example, *Alabama's Board of Corrections* applied the same standard of height to every applicant for a job as a Corrections Officer: at least 5 ft. 2 ins. (*Dothard v. Rawlinson*) This facially neutral criterion excluded a much higher percentage of female candidates than male, and the defendant was required to show that 5 ft. 2 ins. height was needed for the job. The defendant submitted that strength was needed for the job, and height was an indication of strength. The US Supreme Court rejected as unargued the claim that strength was needed, and ruled that if strength was required it should be measured directly. (Levin 2018, 573-574). It is also clear from case law that the employer showing that the basis for the employment decision was a job-related criterion consistent with business interest can be a defense against a discrimination suit. There are also exceptional situations where sex, religion, or national origin (but not color or race) may be used as a basis for an employment decision, but a necessary condition for this is that the trait be a bona fide occupational qualification (BFOQ) or reasonably necessary to the normal operation of the particular business. (Friedman 2017, 35, 94-102) So, even though qualifications are not the only legal grounds for employment decisions, they play a central role in the evaluation of practices, and proved use of them can be sufficient, and is sometimes necessary, for a defense against a discrimination claim.

filter example, and in the law, on this definition of fairness no malicious intent or motivation is required to make a decision unfair.

The conditional in the *Fairness* condition is subjunctive, a counterfactual conditional¹⁹ that puts a constraint on the distribution of jobs and identities not just in the actual world but in all close possible worlds, worlds that aren't actual but might have been. There are several reasons why the conditional needs to thus have modality. Consider first a scenario where the antecedent of the conditional is actually false for some a and b in the applicant population due to their having different levels of qualification. If the *Fairness* condition were a truth-functional material conditional, then that false antecedent would be enough to make it true for a and b, but that is clearly not enough for fairness. For that we also want that if they *had* had the same qualifications they would have had the same chance at the job. Recall the four filtered applicants above and suppose one of them was less qualified than any other applicant in the pool. The antecedent of *Fairness* is false for him as x and anyone else as y, and his outcome of not getting the job even fits his qualifications, but this applicant was nevertheless discriminated against by the filter.

It is sufficient for a material conditional to be true if both the antecedent and the consequent are actually true, but again this is not enough for fairness. Suppose our four filtered applicants were all unqualified, and all of the unqualified *unfiltered* applicants also failed to get the job. Then all in this qualification group got the same (fitting) outcome – not being hired – but the filtered applicants were still discriminated against, according to intuition and the law.²⁰ Being subjunctive, the *Fairness* condition captures the idea that a procedure can be unfair even if no one actually is denied a job they were qualified for.

The subjunctive conditional is also needed in order for the concept of fairness to support the possibility of progress over time. Before the first member of an identity group is actually hired in a given position it is possible for there to be no members of that group qualified for the position, for example because of unavailability of training due to their identity. In such a case a material-conditional version of *Fairness* would count it as fair if an employer had an explicit policy of not hiring people of that identity group at that position. Since in the described scenario no actual member of that identity group had the same qualifications as anyone who was hired in that position²¹, the antecedent of *Fairness* was false, and a material-conditional version of the condition thereby automatically true. A fair policy would instead make us ready

¹⁹ The term “counterfactual” literally refers to a conditional whose antecedent is presumed by the speaker to be false in the actual world. A subjunctive conditional, such as *Fairness*, does not carry that presumption, but the term “counterfactual” has come to be used as a more general heading covering both kinds of conditional. (Starr 2019)

²⁰ See footnote 14.

²¹ Note that for the reasons discussed in this paragraph it would not fulfill *fairness* if we used a sex filter even if it was based on a frequency statistic that said *all* of the actual orthopedic surgeons are male, or all of the pediatricians female.

to recognize the first-ever qualified candidate in an identity group by always taking into account its possibility. The *Fairness* condition must be modal to insure this.

More broadly, that equally qualified candidates get the same outcome (or chance at it) should not be accidental, and making *Fairness* a condition over all close possible worlds adds such an assurance. I take fairness to be a property found in a procedure or method, its having a disposition to produce results that are the same for all possible members of a qualification class. Understanding *Fairness* as articulating a property of procedures, that ranges over possibilities not yet seen, also means that the property can give applicants prospective assurance that they will be treated the same as their qualification-twins.²²

Fairness is necessary but not sufficient for a decision procedure to be just, as I will understand the concept of justice here. Notice that I will satisfy *Fairness* if my procedure is to deny all applicants a job, regardless of qualifications, and give them all a booby prize as I send them out the door. If I treat all applicants the same then a fortiori I treat all applicants with the same qualifications the same, and what I'd done wouldn't have been illegal, but, assuming that some candidates were qualified, some of the outcomes would not be fitting. Analogously, if our policy was not to put criminal defendants on trial, but rather to simply put every defendant in prison for five years, then we wouldn't be treating the set of defendants unfairly relative to each other by the *Fairness* condition, taking "q" to be a binary that separates defendants from the rest of the population. But this procedure would be unjust; it would treat every defendant the same, but would come to ill-fitting outcomes where the innocent were punished, and some guilty people might not be punished enough.

Fairness says that nothing other than qualifications and circumstances²³ may determine the outcome for a candidate, but it does not tell us in what way a given set of qualifications should affect the decision of whether to hire that person. To get at this idea we need to pair getting a job (promotion, etc.) with being more qualified, and not getting a job (promotion, etc.) with being less qualified, rather than the other ways around. Realistically, there are grades of qualification, and sometimes even incommensurable qualifications that cannot be arranged on a single scale. But for current purposes I will make the simplifying assumptions that, for a given job, qualifications can be represented on a single scale, and that there are only two grades: qualified and not qualified. Realistically, competition for jobs is often a zero-sum game, where there are more applicants than jobs, but for simplicity, for the current point, I will assume that there are as many of the jobs in question as there are qualified applicants for it; everyone could have such a job if he were qualified.

²² It is well known to be difficult to give general criteria for deciding which possible worlds alternative to the actual world count as similar enough to be within the scope of the evaluation (Starr 2019). This is a problem that attends any use of counterfactuals to analyze concepts, and I don't have a novel solution to offer.

²³ Generic facts about the world and how it works will also affect the evaluation. They will be part of what defines the set of possible worlds that are close to the actual one, which are the only possible worlds considered in the evaluation of the counterfactual.

What we need to add to *Fairness* is a requirement that hiring decisions fit the applicants' qualifications. Putting (a special case of) *Fairness* together with a requirement of fit, and making it explicit that the decision procedure is the bearer of the property, we have a *Just Procedure* if and only if for all x , and the circumstances, S , and qualifications, Q , relevant to the job, J :

$$Q(x).S(x) \rightarrow J(x) \quad (1)$$

$$\text{not-}Q(x).S(x) \rightarrow \text{not-}J(x) \quad (2) \quad \textit{Just Procedure}$$

The procedure is such that in circumstances S , if x were qualified then he would get a job, and if x were not qualified then he would not get a job. A just procedure has features that make it systematically choose which people to give jobs on the basis of their qualifications, and none of their other features. *Just Procedure* implies a special case of *Fairness*; we're now assuming Q and not- Q are the only levels of qualification, and if $Q(a)$, $Q(b)$, $S(a)$, and $S(b)$ were all to hold, then, by *Just Procedure* both a and b would get the job. Similarly for not- $Q(a)$, not- $Q(b)$, etc., and not getting the job. Unlike *Fairness*, *Just Procedure* requires not only that a and b with the same qualification status get the same result, but that the result fits the qualifications in a preferred way.²⁴ As with *Fairness*, violation of *Just Procedure* does not require malicious intent or motivation.²⁵

To see how it works, consider the sex-filter example above. Suppose Artem is a man who's been filtered out of the pediatrics job, and suppose he was not qualified for the job anyway. A discriminatory filter was part of the basis for his not getting the job, and the law will hold the employer liable. *Just Procedure* ought to deliver the same verdict. Notice that, assuming that the selection process after the sex filter was used depended only on qualifications, it was not the case that Artem would have got the job but for his sex; the sex filter was not the controlling cause of his fate. His not getting the job doesn't personally provide a witness against (1) because he doesn't fulfill its antecedent. He doesn't provide a witness against (2) because though he fulfills the antecedent, he also fulfills the consequent. Where, then, is the discrimination?

The discrimination comes from a property of the filter, that it makes (1) false for any possible qualified male candidate for the pediatrics position, and any possible qualified female candidate for the orthopedics position. A *Just Procedure* is required to support counterfactual

²⁴ As a result, although deciding by a coin toss would satisfy *Fairness*, it would not satisfy *Just Procedure*.

²⁵ One might worry that the definition of *Just Procedure* is susceptible to manipulation. Could we not say that our procedure is to give jobs to all and only applicants who are well-qualified – i.e., $(Q \text{ and } S) \rightarrow J$, and $(\text{not-}Q \text{ and } S) \rightarrow \text{not-}J$ – and mistakes that lead to unfitting outcomes are all errors of application? Our decisions would then always be just, because they fulfill the counterfactuals of *Just Procedure*. However what is part of one's procedure and what is application is not a matter for stipulation. A mistake of application is a one-off – perhaps one lost three pages of someone's CV. Any aspect of the decision-making that is repeatedly applied or followed as a regularity is something you would also do for possible candidates who are not actually being judged. As such, what you would do falls within the purview of what the counterfactuals put conditions on. The conditions do not stipulate what the procedure is, but any part of what you do that supports counterfactual outcomes falls under the scrutiny of the *Just Procedure* conditions. This is the sense in which those conditions are about procedures.

(1) for all possible candidates in close worlds, and the filter procedure guarantees that it will be false for some. My account thus implies that the counterfactual in Fricker's definition – the speaker has not been accorded the credibility (here, job) she would have gotten had she not been judged by means of an identity-generalization – is not a necessary condition for injustice. It is not necessary for a chosen individual to lose something he would have had in order for his treatment to have been unjust. The injustice resides in the way the decision was made, which we can see by considering its consequences for all possible applicants.

On this view a discriminatory hiring decision is morally wrong; it may be accompanied and encouraged by, but does not depend on, and is definitely not the same as, an epistemic mistake. The moral wrongs are unfairness and use of an unjust procedure, one that guarantees that some possible applicants would get an unfitting outcome on the basis of identity traits. This view makes sense of our calling the wrong of employment discrimination an injustice, even if these concepts were not involved in the genesis of the law. The view explains in particular the wrong of using even well-evidenced statistical correlations between identities and expertise in moral terms. You may have reason to believe the correlation in the actual world is as high as 100%, and it may actually be 100%. As long as all that you know is that it is a correlation, it is a moral wrong to use the filter because doing so guarantees you won't hire possible exceptions.²⁶ It would also be to treat candidates differently from their qualification twins, and that is unfair.

I have developed the definition of *Just Procedure* out of intuitions about fairness, being treated the same as others who have the same level of the property that is relevant to the decision, here qualifications, and the fit of job offers to qualifications rather than to their lack. Whether one has been treated fairly depends not only on the outcome but on the rule under which that outcome was determined; rules are general and there is a certain necessity to the relations they impose. Thus the conditions are expressed counterfactually, because counterfactuals put conditions on what would happen in all of the alternatives that the rule might have been applied to. However, we could explain the conditions of *Just Procedure* and why they are counterfactual in another way: what matters to whether a hiring decision was just is the *basis* on which it was made. The proper basis for hiring decisions is qualifications, and the basing relation is counterfactual. Thus for a job offer to be based on qualifications is for it to be the case that if a person were qualified they would be offered the job, and if they were not qualified they wouldn't.

The difference between being judged by the content of our character and by the color of our skin is not the difference between getting and not getting a job offer. It is the difference between the decision to offer a job or not being made on the basis of, say, qualifications and on the basis of identity traits. The notion of something's being the basis, or part of the basis, of a decision is clear enough to common sense, but the concept needed here hasn't been given a

²⁶ If it is not merely a correlation, then there is some kind of law-like connection between the identity and the qualification, and this may be evidence against the possibility of exceptions. Worlds in which white Americans were experts on the experience of African Americans would not be close to the actual world.

philosophical analysis that I know of. The closest topic would seem to be the concept of a reason for action, a justifying or explanatory or motivating reason, but philosophical discussion of these sorts of reasons for action primarily revolves around the question what sort of thing a reason is, whether psychological or objective (Alvarez 2017), whereas what is needed here is an account of what relation between a decision and another thing makes the latter the basis of the former.

The discussion of the basing relation in epistemology is more relevant because there the question is about the relation of a belief to its cause or to the subject's reasons. (Korcus 2019) However, both this and the ethical discussion largely assume that a subject knows what the reason or basis for her action or belief is, or even has 1st-person authority over what it is, and these things can't generally be assumed in the current context for a number of reasons. In discrimination shown by disparate impact of a facially neutral criterion, a conscious purpose of disadvantaging members of a group is not necessary. In disparate treatment cases of discrimination, even if an employer denies it intent can be inferred, from, for example, proof that his explanation is a pretext, or from the existence of a comparator who got a different outcome despite being sufficiently similar in every way except identity. (Friedman 2017, 6-7; Sullivan 2009) It is clear that disparate treatment is not taken by the Court to be limited to conscious, purposeful actions, as shown, for example by its regarding decisions affected by stereotyping as disparate treatment.²⁷ And some kinds of explanatory and motivational reasons are taken to be irrelevant to whether a decision was discriminatory.²⁸ Not choosing a pregnant woman for an assignment out of a desire to protect her is discrimination. The administrator's purpose in using the sex filter was efficiency, but using the sex filter was discrimination.

Abstracting away from what sort of thing the bases of a decision must be, the whole basis of a decision is the totality of things that make the decision go one way rather than another, and this is why at a first pass simple counterfactuals give the best analysis. B is the basis of D when if B hadn't been the case D wouldn't have been the case and if B had then D would; B is the "doer" in D's happening. The whole basis of a decision will typically have multiple factors of various kinds, each making a difference to how the decision went. By Title VII a decision that is said by the courts to be "based on" or made "on the basis of" or "because of" one of the protected identities need only be partially based on an identity characteristic. It is "relied on ... in making [the] decision", "used ... in making [the] employment decision" (Macy v. Holder 2012), not necessarily used exclusively. (Sullivan 2009, 199; Friedman 2017, 15-16) What it is to

²⁷ Stereotyping, even when unconscious, is considered a form of intentional discrimination: "[W]e are beyond the day when an employer could evaluate employees by assuming or insisting that they matched the stereotype associated with their group, for [i]n forbidding employers to discriminate against individuals because of their sex, Congress intended to strike at the entire spectrum of disparate treatment of men and women resulting from sex stereotypes." (Price Waterhouse v. Hopkins 1989, 251) See Eisenstadt and Boles 2016 for discussion of the meaning of intent in discrimination law.

²⁸ "Title VII prohibits discrimination based on sex whether motivated by hostility, by a desire to protect people..., by gender stereotypes, or by the desire to accommodate other people's prejudices or discomfort." (Macy v. Holder 2012)

qualify as a partial basis for a decision could be hard to define neatly, and my positive approach here setting out what the whole basis of the decision should be as one simple thing – qualifications – avoids that complexity; anything that makes a difference to the decision that isn't a qualification is going to undermine one or the other of the *Just Procedure* counterfactuals. So a decision that relies on identity to any extent gets classified as impermissible.²⁹

A just procedure will base decisions on the proper thing for that type of decision – for employment decisions it is (ideally) qualifications – and the basing relation is counterfactual. The definition given above delivers fairness automatically because it implies that everyone with the same complement of the proper thing would get the same outcome. The counterfactuals allow that fair treatment doesn't only depend on actual qualifications or outcomes. Having a condition over all possible candidates promotes progress over time by making us ready to recognize the first qualified candidate of an identity group, and it gives prospective assurances of fair treatment to candidates. The conditions of *Just Procedure* explain why use of identity characteristics, even characteristics highly correlated with job qualifications, is an injustice. A better explanation of the injustice of using identity in decisions will have to do better in these respects.

What does a job applicant deserve? He deserves to be treated fairly and judged by the properties he has that are relevant to doing the job, that is, his qualifications. He deserves his qualifications to be the basis of his getting the job or not, and the basing relation is given by the counterfactual conditionals of *Just Procedure*. Title VII law contributes to realizing this ideal by prohibiting some of our most likely violations of it.

IV. Epistemic Justice: What does a speaker deserve?

There is a strong analogy between employment discrimination and testimonial injustice. If four students raise their hands to make comments in a class, I may commit a testimonial injustice if I call on the lone man of the group first. Arguably I won't be if the reason I call on him is that I know he's written a thesis on the topic, and I know the others haven't. However, I definitely will be if the reason I call on him is that he's a man. What matters is the basis of the decision. If I call on him because he is a man, that implies that if he weren't a man but were equally qualified to speak on the topic I would not, or at least might not, call on him, and if he were a man without qualification to speak then I would call on him, violating *Fairness* and both clauses of *Just Procedure*. If I didn't know that this man had written a thesis on the topic, but called on him because the topic was something mechanical and I knew that only 10% of mechanical engineering bachelor's degrees were awarded to women (Carnevale et al. 2011, 14), I would

²⁹ This will be so except in rare cases where an identity is part of the qualifications for a job. To re-iterate, though it is a useful heuristic this positive approach is not consistent with Title VII law. See fn. 16 for both of these points.

likewise be violating these requirements. By this procedure, if one of the women had the same degree she nevertheless would not be called on.

Though such common examples fit the mold, it may seem that the conditions above inspired by employment discrimination law are not suited to an analysis of testimonial justice and injustice, for three reasons. One is that our processes of deciding whether and who to believe are intuitive, without the rigor of these conditions. Another is that it isn't always appropriate to judge peoples' qualifications to speak. You might think, and some philosophers have argued, that we are ordinarily defeasibly justified in believing speakers by default, without making inferences about their reliability. (Coady 1992, Reid 1764, Burge 1992) (Others regard this as gullibility. (See Fricker, E. 1995.)) Thirdly, one might think that applying for employment and promotion is competitive in a way that testimony often is not.

To the first objection, rigor is not the crucial aspect of the conditions defined above. What is crucial is the propensity and systematicity of our decisions to trust or not, and this is very much present in both our intuitive judgments of speakers and our intuitive judgments of which such decisions are just and unjust. When Herbert Greenleaf dismisses Marge Sherwood's suspicions about the talented Mr. Ripley, saying 'Marge, there's female intuition, and then there are facts', (Minghella 1999) his decision to dismiss her is based on a trait – being female – that would make him be to some extent dismissive toward any woman in similar circumstances, namely where she was casting doubt on someone he favors. If he weren't following such a rule but responding only to some individual traits he distrusts in Marge, then he wouldn't have invoked "female intuition" and it also wouldn't seem like an obvious case of injustice. To respond as a rule the same way to people who share a trait is all that it takes to be following a procedure in the sense discussed above, because it supports counterfactual claims about how the hearer would respond if a speaker did and did not have that trait.

The second concern is the possibility that this framework is inapplicable to large swathes of our decisions whether to trust a speaker, because, especially in everyday life, these are not based on inferences and judgments of speakers' qualifications, but are automatic and defeasibly justified by default. These claims about justification and the process of coming to trust a speaker may be granted or not without affecting my account. That we have a defeasible entitlement to believe others' testimony could be expressed in the current idiom by saying that there are many speaking jobs for which everyone is *prima facie* qualified in virtue of being human. What *Fairness* requires is only that everyone with the same qualifications be treated the same, and using a *Just Procedure* includes this plus the requirement that the treatment go in the appropriate direction. It doesn't matter whether it is done with inferences or without, as long as it is done. A procedure need not even involve a belief. I discussed employers' decisions above without making reference to beliefs about qualifications. Moreover, my account would be able to identify it as unjust when on the basis of identity traits we don't grant the default credibility that we should to someone on a topic where everyone is *prima facie* qualified.

In the actual world, trust by default does not seem to be universal or always applied fairly among human beings, and at least some everyday choices of whom to trust are made by unconscious use of stereotypes. All of those choices made using stereotypes would fall under the purview of the conditions for fairness and justice, as they are currently stated; they are systematic and differential ways of responding to people. Also, situations are not so rare in which we want to learn something that not every member of the human population would know, so we have to judge who is best qualified. Since our attention is a limited commodity, and speakers sometimes disagree, the choice of which speaker to listen to also can be much like having one job to hand out and four people who apply for it.

Though such cases exist, what I am building is nevertheless not a competitive model of the just distribution of a finite good. (Fricker, M. 2007, 20) No zero-sum game is required in order to describe a just procedure as defined above, and there need be no competition between people, just as there may not be for jobs. *Fairness* and *Just Procedure* do require us to make comparisons, but not between actual competitors. The comparisons are between how the person is judged and how other possible people with his qualifications would be judged, whether they are actual or not. The fate of actual competitors can serve as evidence of an unjust procedure, but is not part of what the injustice is.³⁰ Another respect in which my account is not a model of the distribution of a good is that whether an injustice has occurred is not determined by outcomes, or by the fit of consequences to qualifications; Artem above got a fitting outcome, but his treatment was unjust because the decision was made on the wrong basis.

In the speaker-hearer setting, the analog of job qualifications is reliability, which includes both competence – possession of the knowledge – and sincerity – truthfulness in assertions. For a given speaker, competence would vary with subject areas and circumstances, and truthfulness might too, so a speaker's reliability will be evaluated relative to subject matter and circumstances, analogous to the way qualifications are evaluated relative to a particular business, job, and other circumstances. Being hired for the job corresponds in the testimonial context to being granted by the hearer an attitude of receptivity, which in a given case might be as much as believing what they say or might be as epistemically non-committal as paying attention.

We can use the *Fairness* and *Just Procedure* conditions in this context by substituting different properties *mutatis mutandis*. With level of reliability, r , circumstances s , receptivity C , and any individuals x and y , a disposition to be receptive or not to a speaker is *fair* if and only if

³⁰ In one method of proof of discrimination, a plaintiff identifies "... a "comparator," another similarly situated worker who was treated better than the plaintiff was treated. By definition, the more alike a putative comparator is to the plaintiff, except for a protected characteristic like race, the fewer nonracial reasons exist to explain a particular decision." (Sullivan 2011, 1660-1661; 2009) The existence of an actual comparator can contribute to proof of discrimination, but is not necessary for proof, or for discrimination. There may be no female applicant for the pediatrics job who has similar qualifications to Artem, but there was still discrimination.

$$r(x) = r(y) \text{ and } s(x) = s(y) \rightarrow (C(x) \text{ and } C(y)) \text{ or } (\text{not-}C(x) \text{ and } \text{not-}C(y)) \text{ or } P(C(x)) = P(C(y))$$

Fairness

If x and y had the same reliability in the same circumstances (for example the topic of discussion, the conversational setting) then x and y would have the same (probability of) receptivity from the hearer. In this simple model receptivity is a binary – either the hearer pays attention/believes or does not – so this means that either x and y would both get attention or both be ignored, or the two would have equal probability of being heeded.³¹

Lack of this kind of fairness is central to many obvious kinds of testimonial injustice. I may not fuss if you pay no attention to what I say, but I may record it in the book of death by a thousand cuts if you do pay attention when a person of a different identity group says the very same thing three minutes later. Our saying the same thing is defeasible indication that I am at least as reliable as the speaker who repeats my point; it would be quite a coincidence if two people said exactly the same thing within earshot of each other and the second was reliable while the first was not. Your nevertheless paying attention when a person of one identity says it and not when a person of a different identity said it is evidence that identity was part of the basis of the decision whether to listen.

As with justice in employment, we also do care that attention or believing what a speaker says follows reliability and disregard or distrust follows unreliability, not the other way around. So for justice we need:

$$R(x) \text{ and } S(x) \rightarrow C(x) \quad (1)$$

$$\text{not-}R(x) \text{ and } S(x) \rightarrow \text{not-}C(x) \quad (2) \quad \textit{Just Procedure}$$

Updated with the binary property of being reliable or not, R, we have that in order for my choice whether to trust a speaker to follow a just procedure it must be that for every x, if x were reliable I would be receptive to him, and if x were not reliable I would not. The basis of my choice whether to be receptive or not is his reliability and the circumstances, and nothing else.³²

This account may seem more suited to the competence than to the truthfulness aspect of reliability in a witness, but it applies the same to both. To see this consider Fricker's example of Tom Robinson, mentioned earlier. A Black man in Alabama in 1935, he stands trial for the charges of raping and beating a young white woman. The circumstantial evidence leaves overwhelming doubt about his guilt – for example, she was hit from the assailant's left and his left arm is paralyzed – but the framework of interpretation and perception among the white people, and so the whole jury, turns everything he says into evidence against him. His version

³¹ This could be generalized beyond a binary C, to include degrees of receptivity, with fairness requiring reliability to make the same direction and degree of contribution to x's and y's outcomes.

³² That is, nothing else except the properties of the actual world, such as natural laws, that determine which set of possible worlds is close, hence relevant to the evaluation.

of events is the most plausible of the witnesses', but the white people aren't receptive to him³³, and he is convicted. The injustice of the event was summed up succinctly ahead of time by Atticus Finch, the defense lawyer, the day he was assigned the case, when his 8-year-old daughter asked whether he would win: "No, honey. ... [W]e were licked a hundred years before we started." (Lee 2010, 84)

Finch knows ahead of time that it won't matter what he says in court. It is clear in this case that because he was a Black man Tom Robinson would be convicted whether he was guilty (lying) or not (telling the truth). In particular if he were truthful he might not, in this case clearly would not, be acquitted, violating (1) of *Just Procedure*. The basis of his conviction was that he was Black. *Fairness* is violated if he might have been acquitted had he been an innocent white man in the same position who said the same things and was prosecuted with the same evidence. It is likely in this milieu that a white man would have been acquitted, partly because of the evidence and that attention would be paid to it, partly because the pecking order in the social hierarchy would have put the alleged victim woman below an accused white man. In such a case she might have ended up shamed as a whore.

Just Procedure sets the bar high for how we should make decisions if we would be just. By this criterion justice does not merely require that we are justified in believing our procedure fulfills the conditions, or that it fulfills them as far as we know, but that it actually does fulfill those conditions. Thus, if I call on the lone male student with his hand up only because I know he has written a thesis on the topic, that's not sufficient for following a just procedure. I would be violating *Fairness*, hence *Just Procedure*, because if one of the women with their hands up had written a thesis that I didn't know about, I wouldn't call on her. What I did – acting on the basis of information that was differentially incomplete over identity characteristics – displayed a disposition to treat equally reliable speakers with different receptivity.

The conditions are strong – implying that any time my information about candidates' relevant qualifications is differentially incomplete, I should choose between them at random – but for a conception of an ideal of fairness and justice, this does not strike me as wrong. If one doesn't want to choose a speaker at random one has the option of gathering more information about the candidates whom one doesn't have the relevant information about. There is an analogous expectation in hiring that we not make decisions on the basis of lopsided acquaintance with the candidates' files.

Satisfying *Fairness* and *Just Procedure* may or may not be sufficient for achieving epistemic justice. This depends on whether we think that all of the moral wrongs identified by Fricker as potentially present in a case of testimonial injustice will be done on the basis of something other than the person's epistemic reliability; if so, then they won't exist if *Fairness* and *Just*

³³I don't say they don't believe him because it is made explicit in the book that the white townspeople don't believe the accusers. (Lee 2010, 275) Whether the jurors believed or not appeared not to matter to what they would do. But their attitudes and verdict are sufficient to take them as being unreceptive to Tom Robinson's testimony.

Procedure are satisfied. Arguably, the profound lack of receptivity that objectification entails has to be based on something other than the person's epistemic reliability, since every human being is reliable about some matters; every human being is a knower. There are rich questions to explore in this area, but I won't pursue them further here.

Fairness and *Just Procedure* can be useful for defining epistemic injustice as long as they can point to conditions that are necessary for justice, hence sufficient for injustice, but I don't think we regard all violations of these conditions as epistemic injustices. The conditions are apt because they mark, correctly I think, that there is something that employment and receptivity decisions should ideally be based on: qualifications and reliability, respectively. However, there are ways for these conditions to fail that we would not police as injustices. In the employment context, an employer may base a hiring or promotion decision on many things besides qualifications and not run afoul of Title VII. She may promote an incompetent person because the person wears blue shirts every day, or even because the person is a family member or lover, for example. Analogously in the speaker-hearer context, I might listen to some people more than others for reasons that are neither their greater reliability nor their identity traits. I might listen to family members, friends, and colleagues more than others, even if I think their beliefs are false, if what they think and say will affect my life and my ability to coordinate with them. Families, friendship circles, and professions can act as excluding in-groups, but this kind of partiality is not primarily what the concept of epistemic injustice was designed to identify.

To use the *Fairness* and *Just Procedure* conditions to define epistemic injustice, we need to identify a relevant subset of types of failure of the conditions. A natural candidate for this is the set of failures where a person wasn't given receptivity, and might not have been given it had they been reliable, but receptivity might have been given had the person been reliable and had a different legally protected identity trait: race, color, religion, sex (including sexual orientation and gender identity), national origin, disability status, or age. Tom Robinson did not get receptivity for his testimony even though he was truthful, but receptivity would likely have been given if he had been a white man, regardless of whether he had been truthful, which implies the needed weaker claim that receptivity might have been given if he was white and truthful.

Why define epistemic injustice indirectly, via a subset of failures of conditions for justice? Why not define it as one's being denied receptivity on the basis of a protected identity characteristic? The natural way to do the latter would be to make it a condition for injustice that one's identity made a difference to one's outcome, as in Fricker's account above. Assuming a was denied receptivity, we would add a counterfactual with identity characteristics in the antecedent, so:

If a had not been in identity group I, then a would have received more receptivity.³⁴

Identity

I think this condition, on the assumption a was actually denied receptivity, is sufficient for injustice, and it would violate one or the other of the conditions in *Just Procedure*. Either a is reliable or not. If a is reliable then their not actually receiving receptivity indicates a violation of (1) of *Just Procedure*, and the proposed condition *Identity* would tell us this failure was due to their identity. If a is not reliable, then *Identity* would tell us condition (2) of *Just Procedure* is violated, because it tells us if they had had a different identity the procedure used would have given them receptivity despite their unreliability, and would tell us identity was the reason.

The problem with such an approach is that the condition *Identity* is not necessary for epistemic injustice. One's identity may be the basis of a decision – making it unjust – without one's identity making a difference to the outcome. Artem above was not chosen for the pediatrics job because he was a man. This is so despite the fact that he would not have gotten the job had the decision been based on his qualifications, and a person with his qualifications would not have gotten the job if they had been a woman. The outcome would have been the same either way. Given that injustice will not always give the speaker a deficit relative to what they would have gotten without the use of identity, we need something that will identify the reference for what they should have gotten. This is why we need to name something – qualifications, reliability – as the proper basis for a just decision, to give an account of the concept of injustice.

Outcomes are an incomplete guide to the basis on which a decision was made, so we can't take difference in outcome as a test for whether identity was the basis of a decision. In cases where the outcome fits the qualifications it can be difficult to know what the basis of a decision really was, but sometimes there will be clear evidence. In Artem's case use of the sex filter confirms that the basis was sex. Why is the basis of the decision and not a difference in outcome what is important for injustice? Because even if your qualifications mean you're bound to fail, it is unjust for your identity to prevent you from having a chance to compete.

If a person wasn't given receptivity, then I will take it that person has suffered *epistemic injustice* if

- 1) even if they were reliable they might not (would not) have been granted receptivity, and
- 2) receptivity might (would) have been given had the person been reliable and had a different legally protected identity trait: race, color, religion, sex (including sexual orientation and gender identity), national origin, disability status, or age.

³⁴ One may wonder how this differs from *Fairness*, which requires that people with the same qualifications and different identities get the same treatment. *Fairness* includes counterfactuals about all possible people, not just a. So, it requires that Artem and anyone he has the same reliability as get the same receptivity, but it also requires that someone with more reliability than Artem also gets the same receptivity as those with the same higher reliability and a different identity.

Such a person has been subject to a failure of condition (1) of *Just Procedure* – if they were reliable they might not have received receptivity. Clause 2), specifying that the outcome might have been different had their identity been different, identifies the relevant subset of failures of condition (1) as those that turn on (protected) identity. Clauses 1) and 2) imply that *Fairness* has been violated. The “might” in these conditions is sufficient for violation of condition (1) of *Just Procedure*. A higher threshold for epistemic injustice can be had by substituting “would” and “would not” as indicated.³⁵

I have omitted violations of condition (2) of *Just Procedure* as sufficient for epistemic injustice in order to follow the standard restriction of this classification to identity-prejudicial credibility *deficits*, but one could include it if desired. If condition (2) is violated (identity-wise) by a procedure then someone might receive receptivity on the basis of their identity despite being unreliable; the procedure would be classified unjust because it allows credibility excesses. Though it seems infelicitous to say that a person who has been treated favorably on the basis of an identity characteristic has been done an injustice, it seems natural to say that they have received unfair treatment. Even if a decider has committed an injustice in such a case it may seem questionable whether she has done a wrong *to* the person on the receiving end, but that may be because we assume they haven’t been harmed, and a person may be wronged without being harmed.³⁶ This question presents a rich set of issues that I will not try to settle here.

Even without adding condition (2), though, some cases of credibility excess will count as injustices by my criterion, namely those in zero-sum situations. If there is only so much receptivity to go around, then a procedure that allows a positive credibility decision on the basis of identity for one party is one that would allow a negative credibility decision on the basis of identity for someone else, violating condition (1) of *Just Procedure*. This implies that certain kinds of corrective practices will count as unjust. For example, a policy of always upgrading one’s credence in what women say would be unjust in zero-sum contexts because it would result in automatically downgrading credence in what men say, regardless of whether they are reliable. This is an extreme policy, and the verdict of condition (1) seems to me right. But we don’t live in an ideal world, and there may be more targeted policies that seem advisable even if they do violate condition (1). The possible listening policies and the question whether it is wise to permit them might be as complex as US federal and state laws regarding the permissibility of various kinds of affirmative action.³⁷

An example of positive action towards those disadvantaged by their identities arises in Fricker’s conception of testimonial justice, on which this takes the form of a virtue possessed by the

³⁵ Raising the threshold with these “would” constructions will be necessary if we wish to avoid counting a random choice as unjust. See section VI for an example.

³⁶ See also Hazlett 2020, who argues that a credibility excess can both harm and wrong a person.

³⁷ See Friedman 2017, 220-232. Little of a general nature can be said in this area, partly because the need to balance remediation with Title VII’s anti-discrimination clause and constitutional guarantees of equal protection leads to great complexity, and because some states have restrictions on affirmative action that are more stringent than federal law.

hearer. A person with this virtue has a “reflexive critical awareness” that leads them to “revise the credibility upwards” when they suspect the low credibility they attributed has been based on identity prejudice. (Fricker 2007, 89-91) It may seem that in zero sum situations such a person violates condition (1) of *Just Procedure*, but we could understand the described action differently, not as a higher credibility judgment based on identity, but as the hearer basing an upward revision of the credibility judgment on evidence that they’ve made a mistake in the particular case. The “guiding idea” of the practice Fricker describes is to adjust “upwards to reach the degree of credibility that would have been given were it not for the prejudice” (pp. 90-91), which is not an endorsement of upgrading regardless of reliability. The procedure involves paying attention to evidence in a particular case, and the aim of the correction the virtuous person does here is removal of identity as a factor.

Fricker’s virtue-based approach and my procedural approach to epistemic justice are not incompatible in aims, though the latter is more demanding in some ways, but the approach through justness of a procedure is explanatory in ways that the virtue approach is not. No doubt there is a kind of personal moral inadequacy underlying our resistance to evidence that goes against deep-seated stereotypes, and improving ourselves morally at that level can help us, and may even be necessary for us, to perform better epistemically, and so engage in the necessary corrections of our credibility judgments. But this only explains how there is a moral wrong on the part of the hearer, not why what is done to the speaker is a moral wrong, rather than an epistemic mistake, and not why either wrong should be called *injustice*. On my account a person has been done an injustice because they have been treated unfairly, by having a decision about their opportunities be made on the basis of something other than their qualifications or reliability.

My sufficient condition for epistemic injustice uses the legally protected identity traits – race, color, religion, sex (including sexual orientation and gender identity), national origin, disability status, and age – to limit the space of unjust bases for receptivity decisions. Most if not all of these would seem to qualify by Fricker’s standards as social types that identity prejudice can be based on. Arguably they all fall into her special case of *systematic* injustice (Fricker 2007, 27), because prejudices based on these elements of identity track subjects through many dimensions of social activity, and lead to injustices across many dimensions of their lives: educational, professional, legal, medical, financial, and political, for example. This is plausibly part of the explanation of our having laws about these identity characteristics, and not, for example, about the social types *used-car salesman* and *family doctor*. Like the law, and unlike some philosophers, I take the identities whose use in a decision yields an injustice to be categories that include many types, so that a decision is based on race if it is based on any racial type, including white, and not only on those races that have suffered disadvantage as a group. Thus, on my criterion, analogously to the law, it is possible to do an epistemic injustice to a young, white, Caucasian, non-disabled, Protestant, straight, American cis-man.

One reason Fricker limited epistemic injustice to those cases where the generalization behind a decision was unreliable was in order to avoid the definition being over-inclusive. My criterion would be too inclusive if using the identity categories I have listed was a part of everyday life that we could not do without, but I think it is not. Fricker describes the typical hearer as perceiving the speaker "... in light of a set of background assumptions about how far people like him are trustworthy about things like this in relation to people like her". (Fricker 2007, 36) We can accept this picture, and her thought that reliable stereotypes have much to offer, without losing much by refraining from substituting "Black men" for "people like him". There are reliable generalizations about family doctors (dependable), people wearing car-mechanic uniforms (knowledgeable about engines) and, perhaps, used car salesmen (dishonest), but I doubt that there are many reliable generalizations of the form "Black men are/are not trustworthy about things like this". Where there are reliable generalizations involving protected identities and expertise – "Women are not orthopedic surgeons", "Black people are not airline pilots" – I doubt their usefulness in everyday life. Almost no one is an orthopedic surgeon or an airline pilot, so if you want to find one you'll need to ask questions that go beyond sex and race in any case; using sex and race will often be a waste of time. What prevents my criterion for epistemic injustice from including too much is not a restriction to unreliable generalizations, but a restriction on the identity types.

What does a speaker deserve? Sometimes a speaker deserves to be listened to, but always a speaker deserves to be treated fairly. Many speakers deserve more than we will give them, because so much of our limited attention goes to people with whom we have pre-existing relationships not based on judgments of reliability. A minimum to strive for is to avoid basing a decision of whether to be receptive to a speaker on the protected identity characteristics, an idea I have tried to capture above in the sufficient condition for epistemic injustice.

V. Frequency Statistics and Epistemic Obligations

If we're to be just, then we must not use the correlation between sex and orthopedic surgeons to decide whether to give an applicant a job, or use the low percentage of mechanical engineering degrees awarded to women to decide who to call on in class. Still we believe, and even know, many such correlations, and there are credible sources for learning more of them. As things we know, these correlations would be counted by most epistemologists as part of our evidence.³⁸ By ignoring this evidence are we not shirking our epistemic obligations? If I know that the probability that the female student has done a degree in mechanical engineering is 10%, how can I epistemically justify ignoring that? And is there not a trade-off here, where fulfilling the conditions for justice is going to deprive me of knowledge?

³⁸ Though Williamson's (2000) E = K thesis is controversial, most criticism attends the left-to-right direction.

The Requirement of Total Evidence (RTE) – roughly, take all of your evidence into account, not just some of it – is now taken to be so obvious that it is rarely defended, but it is instructive to recall past efforts to justify it in order to identify its proper scope. I will argue that the only cogent defense of the principle does not apply to evidence that takes the form of frequency statistics. If so, then as far as this principle is concerned we are in the epistemic clear if we ignore mere correlations when judging job candidates or speakers.

A version of RTE first appeared in 20th century philosophy in 1934 with Hans Reichenbach's *Requirement of Maximum Specificity*.³⁹ (Reichenbach 1949, 372-378.) This principle is a response to the fact that the probability that an individual has a property depends on the set of evidence that is taken into account. There is not just one probability that Pederson will die next year; there is the probability that he will die next year given that he is a man, the probability given that he is a man, a Swede, and 55 years old, and so on. None of the corresponding probability statements is more true than another, or more relevant to Pederson, but they do not in general have the same values. Which one is the right one to use to judge Pederson, the individual before you?⁴⁰

The Requirement of Maximum Specificity says to choose the most specific classification for which you have reliable statistics.⁴¹ Assuming there is a unique such class, its probability will take into account all of the properties you have evidence about from the other statistics.⁴² Using it will be effectively taking all of the evidence you have into account, since this probability renders those attached to wider, less specific, reference classes for which you have statistics irrelevant; they wouldn't make a difference if you did include them. Though there is a probability that Pederson will die next year given any length of list of properties that he has, you don't know what all of those probabilities are. You can't take into account evidence (probabilities) that you don't have, and the most specific evidence you have – your total (reliable) evidence – is the best you can do.

Intuitive as it is, justifying this choice is difficult. Reichenbach says in its defense that if we follow such a policy we will be right, e.g. about whether given individuals will die in the next year, in a greater proportion of cases. (Reichenbach 1949, 372, 373, 375) But even if this is true it says nothing at all about whether we will be right in the individual case of Pederson, and that was what choosing the right set of evidence was supposed to give us some purchase on.

³⁹ Versions of the principle are found in J.M. Keynes, and before him Jacob Bernoulli, and C.S. Peirce. References to these works can be found in Carnap (1950, 212).

⁴⁰The reference class problem is often taken to be peculiar to the frequency interpretation of probability, but it is not. See Hájek 2007, and Ayer 1957. So I have not used this name for the problem in describing it here.

⁴¹ "Reliable statistics" here refers to your data strongly supporting the particular value assigned to the probability, which could be any value between 0 and 1 inclusive. "Reliable generalization" in the discussion of stereotypes above refers to a generalization that is mostly true; expressed in probability, the probability of one trait given the other is *high*.

⁴² If you have reliable statistics for two equally specific reference classes but not for their intersection (the conjunction of the associated properties), then you have another choice to make. That problem is not relevant here.

Reichenbach points out that this policy will lead to better results for Pederson “as long as the probability is increased with each step” of additional evidence, and that experience teaches us that as we add more and more properties of Pederson to our evidence eventually the probability settles down to a limit, so no further evidence makes a difference. (Reichenbach 1949, 375-376) However we don’t know how long that eventuality will take, and given our limited abilities at evidence-gathering Pederson’s death may come before we get there. We don’t in general know when that limit will be reached⁴³ and it is characteristic of inductive evidence that before that stabilizing limit is reached more evidence may, and often does, make the probability go *down*. When we add the evidence that Pederson is a smoker, the probability of his death next year will go up, but if we subsequently add that he consumes broccoli and red wine each day (but not too much), his probability of death will go down. These points cannot justify the choice of the maximally specific reference class, and so the total evidence, for judging an individual.

The intuition that more evidence is better than less is probably part of what makes the requirement of total evidence seem obvious. But “more is better” is simply false for inductive evidence, because each next bit of evidence may be misleading, that is, take you further from the truth than you were before. The evidence that all the swans you’ve seen were white is misleading in this way. It is subject to a cross-induction by the evidence that there is often color variation between geographically separated populations within a species. No one would blame you for concluding that all swans are white if you have no further information than your observations of white swans, but your confidence that all swans are white should go down when you get the further evidence about species and color. For any piece of evidence that you have, you might later find that for the purposes of getting at the truth it would have been better to ignore it. As a matter of fact, more evidence is not necessarily better, and bigger subsets of your evidence are not necessarily better than smaller ones for getting nearer the truth.⁴⁴

In proving theorems it is often assumed that the evidence stream is not misleading, but in the kind of reasoning we are concerned with here, in application to individuals, in judging their qualifiedness for a job or their reliability as a witness, this is not something we can assume. You don’t in general know of your pieces of evidence whether they are or are not misleading, because for any piece you could later find that it was. So from your point of view all of the evidence that you do have is equally liable (or not) to be misleading. Another idea, then, is that

⁴³ The Law of Large Numbers may be part of what drives an intuition that more evidence is better, but this theorem that says more evidence will make your probability value converge in a predictable way has conditions that are rarely present in casually encountered cases. These conditions are that the trials are independent and identically distributed. Each trial, say the toss of a coin, must have a result independent of the other trials, and the random variables for each trial must have the same probability distribution.

⁴⁴ Note that our question here is not why we should look for new evidence, but why we should let everything (relevant) that we know be the basis of a belief, and not just use some of what we know. The question of the value of looking for new evidence, or for looking at evidence that you have but haven’t yet looked at, is discussed below.

it would be arbitrary, or introduce bias, for you to use some of that evidence and not all of it. This, and perhaps the fact that the choices could be motivated by wishful thinking about the conclusion, is likely the idea behind the examples Rudolf Carnap used to defend his version of the requirement of total evidence:⁴⁵

If a judge in determining the probability of a defendant's guilt were to disregard some relevant facts brought to his knowledge; ... then everybody would regard such a procedure as wrong. (Carnap 1950, 211-212)

Carnap cited Keynes, Bernoulli, and Peirce's endorsement of a total evidence requirement, but offered no further illumination of its justification. But even if taking all of our evidence into account is sufficient to avoid bias, it does not appear to be necessary. We could use our evidence in an unbiased way by using not all of it, but a random sample.

C.G. Hempel endorsed the requirement of total evidence for picking out a preferred reference class, and his analysis of the situation yields insight for justifying the requirement: rationality forbids you believing things that are inconsistent, and if you are allowed to base a belief in p on merely a subset of the relevant evidence concerning p , the rest of the evidence might support not- p , and you would be allowed to also form a belief in not- p on the basis of that. (Hempel 1960, 453-454) It is only if you take all of your evidence into account in forming every belief that you avoid this; all of the evidence will have a tug of war, and between p and not- p only one of them will win.

Patrick Suppes (1966) articulated the most cogent and general version of this justification for the requirement of total evidence when he pointed out that a generalization of consistency, probabilistic coherence – the rationality requirement that says your degrees of belief (credences) should conform to the probability axioms – by itself implies that every degree of belief you have will have taken into account all other degrees of belief that you have in other propositions. This follows from the theorem of Total Probability, which says that for any propositions, p and q , of your language,

$$P(p) = P(p/q)P(q) + P(p/-q)P(-q)$$

The probability of p is equal to the probabilities of p given q and given $-q$, weighted by the probabilities of q and $-q$ respectively. If probabilities are taken as credences (degrees of belief), then your credence in a given p must bear this probabilistic relation to your credence in q for every q . If we require degrees of belief to conform to the probability axioms, then we don't need to add a separate requirement to take all of our evidence into account, or an additional justification for doing so. Any argument that defends probabilistic coherence as a requirement

⁴⁵ "In the application of inductive logic to a given knowledge situation, the total evidence available must be taken as basis for determining the degree of confirmation." (Carnap 1950, 211)

of rationality will be a defense of the requirement of total evidence, RTE, and there are several such arguments.⁴⁶

Thus, for a probabilist about rationality, RTE comes along for free.⁴⁷ This solid defense of the total evidence requirement does not, however, tell us anything about how or whether to take frequency statistics into account.⁴⁸ It provides a constraint on degrees of belief because the probabilities are interpreted as degrees of belief. Any frequency probabilities we might know are not statements of a subject's degrees of belief, but statements of the proportion of one type of event among a class of some typically other type of event. The axioms of probability require that frequency probabilities be coherent with each other, and that credences be coherent with each other, but do not tell us how frequency probabilities and credences should be related to each other. Thus, the best, perhaps only, justification for a requirement to take all of the evidence we have into account is silent on whether, why, how, and under what conditions we may or should take known frequencies into account in our degrees of belief, and so, in particular, in our degrees of belief about individuals. This by itself takes the sting out of the question whether in ignoring a frequency statistic for the sake of justice we are shirking an epistemic obligation; it is not obvious that we are, and it doesn't automatically follow from the ideas that led us to ask the question.

To determine how, when, or whether we are epistemically obligated to take frequencies into account in dealing with individuals we must decide what to think about the following kind of argument, called a statistical syllogism or direct inference:

$P_{FR}(\text{Protestant}/\text{Swede}) = .80$

Swede(Pederson)

$P(\text{Protestant}(\text{Pederson})) = .80$

We assume in this case that all you know about the individual Pederson is that he is a Swede. The first (major) premise is a frequency probability that says that the proportion of Protestants among Swedes is 80%. You might have estimated this from a survey. The second (minor) premise records our knowledge of Pederson, an individual, that he is a Swede. This is not a probability statement but a flat-out assertion that Pederson is Swedish, but even if we replaced

⁴⁶ Dutch Book arguments to the effect that you will avoid sure loss in betting if and only if your credences conform to the probability axioms, are the most well-known type of defense. (Hájek 2008) This style of argument is criticized for being merely pragmatic, but there is also a purely epistemic defense of conformity to the probability axioms, namely, that you will be as accurate as you can be. (Joyce 2009) Representation theorems are another style of argument for defending probabilistic coherence. (Maher 1993, 182-197; Zynda 2000)

⁴⁷ This coherence defense of a requirement of total evidence is not available to a non-Bayesian, which one might see as an argument for Bayesianism.

⁴⁸ On Suppes' view these are taken into account by the frequencies entering our credences with variables, and our conditionalizing on the statement that an individual has a property occurring in those frequencies, and forming a new credence function, but the justification of these steps is not entirely clear. (Suppes 1966, 58-59)

it with a probability statement, $P(\text{Swede}(\text{Pederson})) = 1$, that probability would be a credence and not a frequency, so it could not by logic (or the probability axioms) alone discharge the major premise. It is nevertheless inviting to take these premises to justify having a .80 credence that Pederson is Protestant, as described in the conclusion of the syllogism.⁴⁹

The epistemological significance of the question whether or when we are justified in making this inference cannot be overstated. On one hand the axioms of probability do not compel us or even justify us in taking this step. On the other hand, never taking this step will apparently lead an empiricist to skepticism about induction. As Hume taught us, we can't observe causes. According to empiricism the only evidence we have is what we observe, and this can only be singular facts – such as that Pederson just said he is a Swede – and observed correlations among them. Direct inference is the only way to make use of these things together to form beliefs or credences about cases we haven't seen, so if we are never justified in making such an inference, then we are never justified in using what we have observed to form a belief or credence about what we have not observed. (McGrew 2001)

The significance of this question is attested by the quantity and depth of discussion it has received among philosophers of statistics. I will extract from that discussion just a few points that are relevant here. First, the inference is not permitted at all from the point of view of someone, such as Reichenbach, for whom frequencies are the only probabilities there are. On such a view, probabilities are properties of classes and class relations over populations, and not properties of individuals. So there simply is no probability to be put in the conclusion-place in a direct inference. On this kind of view, either one does not use probabilities to guide decisions about individuals or one adopts a policy, a rule of acceptance, for when the probabilities are right for acting as if the proposition in the conclusion of the syllogism – that Pederson is Protestant – is true. In neither case does the frequency probability affect one's belief, or degree of belief, that Pederson is Protestant.

What one is permitted or obliged to do when one learns frequencies does not affect belief states concerning the individual, but it may affect which actions are deemed choiceworthy. According to a particular rule of acceptance Reichenbach discussed, we would act as if Pederson is Protestant if Swedishness forms the maximally specific reference class and the probability one is Protestant given one is Swedish is $>.5$, as it is in our case. (Reichenbach 1949, 372-378) Using this policy, if all we know is that a reliable estimate of the frequency of orthopedic surgeons among the women in our applicant pool is 4%, as is implied by the labor statistics and the set-up of the sex-filter example, then we would act as if a woman in the applicant pool at the hospital is not an orthopedic surgeon. It would make no sense to consider someone for an orthopedic surgery job who was not an orthopedic surgeon, so we would

⁴⁹ The conclusion of such an argument could also be taken to be a full belief. I am taking it that the inference is one in which the subject ends up with a .80 credence that Pederson is Protestant because credences are, on the Bayesian view, probabilities that force updating of other credences, and credences enter naturally into the expected-utility way of determining the rational course of action for an agent, which I will employ below.

summarily dismiss her file from competition for that job. That particular rule of acceptance would lead us to violate anti-discrimination law, but the conflict is not between an epistemic norm and a practical one; it is within the domain of action. One could try to revise the rule of acceptance to make the actions it leads to conform to the law, but merely raising the threshold higher than .5 won't be enough to do that, since the frequencies we are not supposed to act on can be as high as you like, and often are very high, as in our orthopedic surgery case.

If one thinks that frequencies are not the only probabilities and there are subjective probabilities too, as in the Bayesian paradigm, then the conclusion of the direct inference above makes sense – one's credence that Pederson is a Swede is .80. However the inference is still not automatically justified.⁵⁰ On everyone's view there is a further condition, and even when the condition is met the inference is not forced or obligatory in the way that (for a Bayesian) probabilistic coherence forces one to take one's degrees of belief in all propositions into account on pain of irrationality.

Henry Kyburg (1977) took the most permissive approach to direct inference, and argued that we are entitled to make the inference provided two conditions are met. The first is a condition on the relative precision of the property forming the reference class (which is met in our cases), and the second that for any class we know Pederson to belong to that might be more specific than the class of Catholics among Swedes, we do not know that the frequency of Protestants in that class is different from .80. (Kyburg 1977, 510; 1983, 156) This second condition is that the individual be a random sample, in Kyburg's proprietary sense of this concept; on his conception randomness is epistemic, satisfied by our being ignorant of what effect a further property or its absence has on the frequency of the target property. We might say that for Kyburg, ignorance is bliss for direct inference.

To defeat the direct inference, on this view, we would have to *know* of some property that we *know* our individual to have and that might be more specific, that the frequency of Protestants in that class *is* (not "might be") different from that of Protestants among Swedes. That is, we would have to have reliable statistics about a reference class Pederson belongs to that is different from that of Swedes and might be more specific; the possibility that he is an exception to the frequency statistic makes no difference. The direct inference is legitimate by default, so one would be epistemically justified in forming not only a credence of .80 that Pederson is Protestant, but also for any particular woman in the applicant pool for a job at the hospital a credence .96 that she is not an orthopedic surgeon, before looking inside her file.

Kyburg was in the minority on the conditions for direct inference among writers of his time. On the less permissive end of the spectrum, Isaac Levi put strong requirements on it. (Levi 1977) On his view a frequency does not by itself justify a direct inference even if its condition is the only thing you know about the individual. Such an inference is justified only if, for example in

⁵⁰ Here I am considering whether the resulting credence is justified by the premises, not the question, which I take to be distinct, of whether the inference gives one *knowledge* that Pederson is Protestant. (Cf. Moss 2018, Ch. 5.)

Pederson's case, you also know that the individual has been chosen from the set of Swedes by a process whose chance of picking a Protestant is the same as the frequency of Protestants in the population of Swedes. (Levi 1977, 10) Such a process is random in the sense Levi thinks is needed; it produces a representative sample of the Swedes.

Two ideas lie behind this requirement. One is that the randomness of the choice process transforms a mere frequency probability into a chance, and the other that a statement of chance is what is needed in order to form a credence about whether an individual has the target property. A frequency is a mere correlation in a population while a chance is a probability that represents a law-like generalization that supports counterfactuals. A statement of chance thus says that having the one property contributes to making you have the other one, a relation that would inhere in each individual. The property that has a chance in this case is that of Pederson being chosen by a random process from (being a representative sample of) the set of Swedes. If we know Pederson is chosen randomly in this sense, then on this view we are justified in having an 80% credence that he is Protestant. Otherwise not.

Whereas on Kyburg's view – which I will call *permissive* – more knowledge than the premises offer is needed in order to defeat a direct inference, on Levi's view – which I will call *strict* – more knowledge than the premises offer is needed to justify that inference. Whereas on the permissive view it is enough that we not know that Pederson has a property that makes his probability of being Protestant different from 80%, on the strict view we are required to know that he was chosen in such a way as to screen off the effect of any such property on his chance, as the chosen individual, of being Protestant.

Under Levi's strict condition, a great many inferences leading from frequency statistics to credences about properties of individuals that we might consider making in everyday life are epistemically unjustified. In an example from the literature on profiling, suppose you know that the vast majority of women working in a given office building are administrative assistants, and as you walk down a hallway in this building on your first visit you cross paths with a person about whom you know nothing other than that she is a woman. (Moss 2018, 220; Gardiner 2018, 174) On the permissive view you would be justified in having a high credence that she is an administrative assistant, thanks to your ignorance of her other properties. On a strict view you would have to know a whole lot more, for example about how the circulation of people in the building works, in order to come to that conclusion. A sufficient condition would be that there's a mechanism that makes administrative assistants and executives spend an equal amount of time walking in the hallways, for all hallways. That would insure that the proportion of women in each type of job who will cross a particular point in the hallways is the same as the proportion of women in each type of job. If you know this then you are justified in a high credence that the woman you encounter is an administrative assistant, but even in (presumably rare) cases where this was true, knowledge of it would be hard to come by.

Given the strict criterion's ability to block a lot of everyday profiling, we might then try to epistemically justify ignoring frequency statistics in, say, hiring decisions by opting for Levi's

rather than Kyburg's epistemology of direct inference, or by taking a strict approach when the subject matter is people and a permissive approach to everything else.⁵¹ However, rarely as the strict condition may be met in ordinary living, it is easy to meet it in the context of the sex filter. We meet it when we know of a woman we pick that the chance of this chosen woman being an orthopedic surgeon is the same as the frequency of orthopedic surgeons among women in our applicant pool, namely 4%. Our choice of a given woman has this property if it is done without consulting any information in or of the file except the sex tag, by using a random choice mechanism. Thus at the stage before folders are opened we are entitled, for any woman applicant, to a credence of .96 that she is not an orthopedic surgeon.

No belief state about the individual is epistemically forced on one when one learns frequency statistics before one adopts some epistemological view about direct inference, and even these conditions have largely been understood as giving permission, not obligation. But even the strictest epistemological condition on direct inference that has been contemplated does not always forbid formation of a kind of credence about an individual that may lie behind hiring decisions that violate Title VII. Thus there is a degree of tension between our epistemic and ethical aspirations. One might try to strengthen the epistemic conditions on direct inference further, to make such credences come out as unjustified, but it is hard to see how to do this without inducing a global inductive skepticism.⁵²

At this point a view like that of Rima Basu may seem attractive. (Basu 2019a) On her view, having a high credence formed on the basis of identity statistics in the way we have discussed would usually count as immoral, a belief that wrongs the person the belief is about. A further part of her view is that beliefs that are immoral to have can be thereby epistemically unjustified, which is to say that there is *moral encroachment* on epistemology: the epistemic status of a belief can depend on moral factors. (Basu 2019b, Moss 2019⁵³) This is one way of resolving tension between epistemic and ethical obligations; acts like discriminatory hiring that are unethical to do are based on beliefs that are epistemically unjustified because they are

⁵¹ A dual strategy something like this is used by Moss (2018, Ch. 10) in her theory of knowledge (rather than justified credence that is my concern here). When the subject matter is not persons, she takes direct inference to be successful when no other evidence is present (Moss 2018, Ch. 5), as do some other contemporary authors (White 2010, 169), a condition that is similar to Kyburg's permissive view. But direct inference will rarely give you knowledge about people, on Moss's view, because of the *Rule of Consideration* that she adds to the Lewisian contextualist account of knowledge: "in many situations where you are forming beliefs about a person, you morally should keep in mind the possibility that they might be an exception to statistical generalizations." (Moss 2018, 221) The possibility that they are an exception will be a possibility you must rule out in order to count as knowing they have the property, and a non-extreme statistical generalization will not be sufficient to do that.

⁵² One might try to make the epistemic and ethical conditions match by saying that a direct inference is defeated if you have reason to believe there exist properties such that which version an individual has will likely make a difference to the probability of the target property, though you don't know which version the individual has. This condition hasn't been considered, as far as I know, likely because it would defeat any direct inference to a target property for a phenomenon that wasn't otherwise indeterministic.

⁵³ Note that Moss does not regard the impact of her Rule of Consideration as an instance of moral encroachment. (Moss 2019, 191)

immoral to have. Moral encroachment is unattractive to me, for reasons articulated well by Gardiner (2018), but one might think that such a belief state is morally evaluable and immoral without embracing moral encroachment on epistemology. In that case it would follow that epistemic and ethical obligations are not just in tension but in outright conflict over the (high degree of) belief that, e.g., the woman is not an orthopedic surgeon in cases like that of the sex filter. The credence about the randomly chosen woman would be epistemically justified but morally wrong.

However the demands of justice outlined here do not force us to either of these positions, and we don't need to strengthen the epistemological conditions on direct inference in order for the ethical and the epistemic to be peacefully co-existing sovereign territories. Title VII law doesn't forbid you forming a degree of belief about a candidate on the basis of an identity statistic. It forbids you making a negative *employment decision* on the basis of an identity statistic. As an act, sending a woman's file to the pediatrics department because it is the file of a woman is morally evaluable and unjust, but having a credence of .96 that the woman is not an orthopedic surgeon is not the same as and does not automatically lead to or justify such action on her file. In the next section I will discuss a type of decision that has not been considered that lies between acquiring credences and basing actions upon them, that is morally evaluable and has epistemic and moral consequences. We will see that a second requirement that has run under the banner of total evidence is a friend of just decision-making.

VI. Delayed Decisions – Waiting for Evidence

Belief states and actions are distinct, but how can this be of any use? There is no contradiction in offering a job to someone I believe to be unqualified, but that is hardly helpful since I would also expect this action to have a high cost. Degrees of belief do not improve this situation. A degree of belief doesn't automatically justify an action, but it would seem to with further conditions that are easy to come by: why would anyone offer a job to a person whom they were 96% confident did not have the minimal credential for it?

What we are neglecting is that offering or not offering the job are not our only options in action. If we look at evidence – say a line on Sally's cv that says she had a residency in pediatrics, or a frequency in the type of case just discussed – and we form a degree of belief on its basis, then on pain of incoherence we must update our degrees of belief in all other propositions in light of our new belief state. There is no option to be less hasty in updating your belief states, but you do have a choice about whether to act on the basis of your current belief states and utilities, or wait to act on the basis of the belief states you will have after you see more evidence.

In addition to asking whether we should send file #29 to its demise, we can ask whether we should make an employment or listening decision based on the degrees of belief we get about individuals from the frequencies (in cases where this is epistemically justified), or wait to make

employment and listening decisions until we also look at specific facts in the candidates' files or history. Not waiting for more evidence before deciding whether to send a folder on or call on a student is definitely an action, and as an action is a natural object for moral evaluation.⁵⁴ As we have seen, justice (in the form of *Just Procedure*) weighs in favor of waiting for more evidence because basing hiring or listening decisions on identity frequencies is an unjust procedure. We can also ask whether to act or to wait for more evidence as a pragmatic question, employing a cost-benefit analysis, with justice and better fit of skills to jobs as benefits, and the time and effort of reading files as a cost. Waiting for more evidence before acting in these situations will discourage unjust acts, but it also happens to be recommended by expected utility considerations, via a theorem that has often been called the Principle of Total Evidence (PTE).⁵⁵

In the context of the discussion in the mid-twentieth century of the requirement to take all of the evidence we actually have into account, A.J. Ayer thought that the natural choice among the surfeit of probabilities that an individual event has depending on which evidence is taken into account left it puzzling what reason we would ever have to find new evidence. If the right probability to use is the one attached to the reference class that takes all of your current evidence into account, then what is the value of finding evidence you don't yet have or looking at evidence you possess but haven't yet seen? (Ayer 1957) The statistician I.J. Good answered that question with an argument using the expected utility account of rational decision-making. (Good 1967) Good pointed out that when you have a degree of belief in p , and a decision to make between doing act A and doing act B that depends on your degree of belief in p , you have another choice node: whether to decide on A vs. B using your current degree of belief in p , or wait to decide on whether to do A or B until after you get new evidence that is relevant to p . Getting new evidence (gathering new evidence and looking at it or looking at evidence you have but have not yet looked at) is itself an act, and that act has an expected utility.

Good found something general that can be said about this expected utility, and expressed it in a theorem that says that, on two conditions, getting new evidence will always leave you at least as well off in expected utility as if you hadn't gotten it. Thus the value of new evidence follows from the general expected utility theory of rational choice. The two conditions of the theorem are 1) the cost of gathering and assimilating the new evidence is negligible, and 2) the new evidence could make a difference to which action you choose. Notice that the first condition is not always fulfilled, and if the cost of getting new evidence were always negligible then we should never act, except to get more evidence. In general, costs of gathering new evidence include not only things like the labor involved in finding new data and the computational power

⁵⁴ Notice that while I do think the act of not waiting for more evidence is morally evaluable, and sometimes wrong, it is a distinct act from making a hiring decision on the basis of current evidence. It is the latter, and its counterpart in hearer behavior, that I have argued above wrongs a person by being unjust. Thus my discussion here is consistent with the position I have taken above that faulty handling of evidence is not the source of the injustice in epistemic injustice.

⁵⁵ The phenomenon about which Good proved the theorem that he and others called the Principle of Total Evidence, is sometimes referred to today as the value of knowledge.

required to assimilate them, but also the sheer time spent doing so, and the potential opportunity cost of delaying the decision whether to do act A or act B.

The second condition of the theorem is fulfilled in the cases where the *Just Procedure* conditions above told us not to use the sex filter: evidence we find by looking into a candidate's file rather than only at their sex could make a difference to our credences in such a way as to make a difference to whether we offer them a job. Using the sex filter on the file of a woman who happens to be an orthopedic surgeon would have us deny her the orthopedics job, whereas looking inside her file would have prevented that automatic disqualification because we would know she was an orthopedic surgeon. As for condition 1), in our case the cost of getting new evidence is not negligible: looking at files is labor- and time-intensive. But the cost of getting and dealing with new evidence can be outweighed by other benefits, so that the net cost of gathering new evidence fulfills condition 1). In our case, against the labor and time costs of reading files we have the benefits of a potentially better fit between people hired and job skills, and the use of a just procedure. In the context of employment decisions our society has decided that the latter outweighs the cost in labor and time of reading candidate files.

The value of justice and the Principle of Total Evidence will recommend waiting for more evidence whenever possible in the hiring and listening decisions we have discussed. But in one case we have seen a decision must be made before more evidence is available: I know the man with his hand up in my class has a degree in mechanical engineering and a lot to say, and I must decide who to call on before knowing whether any of the women with their hands up have a degree in mechanical engineering or anything to say.⁵⁶ In the context of a class a question from a non-expert may be just as valuable as a comment from an expert, or moreso, and that would justify calling on a person with their hand up on a random basis. If not, then the greater value of an informed comment over a question might still be outweighed by the greater justice of choosing at random. I needn't postpone calling on someone or other in order to postpone a decision based on my current knowledge (and lack of it) about the students' knowledge.

The PTE will promote justice insofar as (we care about justice and) waiting for more evidence short-circuits a decision based on identity traits alone. But some decisions evaluable for whether they are just or not are themselves assays for more evidence. The decision to interview a candidate is a move to get more evidence before making a job offer, but it is itself also an employment decision by the lights of Title VII, and justice requires that it be made on the right sort of basis, in a way that doesn't disadvantage anyone on the basis of identity. Similarly for a decision to listen to one speaker rather than another. But my claim is not that waiting for or seeking more evidence is what justice is, and we should not be surprised if some ways of getting more evidence are unjust or otherwise morally wrong. Hacking into someone's private account is a way of acquiring more evidence.

⁵⁶ Asking the women whether any of them have a degree in mechanical engineering would be not only weird but offensive – would I have asked men if they had such a degree before calling on them?

The demands of just decision-making do not force us to declare credences based on frequency statistics immoral, or epistemically unjustified. Such belief states may lead some people to unjust decisions, but those decisions will be unjust for reasons that do not depend on the epistemic or moral status of those belief states. As we saw above, such a belief is neither necessary nor sufficient for an unjust decision. A decision is just or unjust depending on its basis and one needn't have credences about job candidates in order to disqualify them on the basis of sex, or any credence about the reliability of a speaker in order to pass them over or brush them aside on the basis of identity.⁵⁷ Nor is a credence sufficient to make a decision unjust, for one may use a formal procedure that doesn't allow one's beliefs about identity and qualifications to have an effect, as in blind grading and blind auditions.

A hammer is a legitimate carpenter's tool, but that doesn't prevent its being used as a murder weapon. Likewise its being used as a murder weapon does not undermine its status as a legitimate carpenter's tool. One might worry that if we don't disapprove of credences about people that are based on identity statistics in the most sweeping and strongest possible terms, then people will not be sufficiently discouraged from unjust decisions. But having a high credence that a woman is not an orthopedic surgeon is no excuse for not looking at a file that for all you know was submitted for the job in orthopedics. Placing sufficient value on justice prevents that credence from underwriting an unjust hiring decision, even on an expected utility analysis. There is no excuse for blindly letting our belief states drive our actions.

VII. Conclusion

The wrong done to a victim of epistemic injustice is not epistemic, and though it may be accompanied and encouraged by a variety epistemic errors, it is not a necessary condition of epistemic injustice that one's choice whether to heed or ignore be mismatched to the speaker's reliability, and not necessary that a generalization one uses to judge an individual be unreliable. It is not necessary for epistemic injustice that a speaker fails to get something he otherwise would have had.

Epistemic justice ideally requires according a speaker credibility or not on the right *basis*, namely his reliability. It implies that a speaker is heeded or ignored to the same extent as any of his reliability twins would be in the same circumstances, and that whether he is heeded or ignored depends only on his reliability or unreliability, respectively. When a decision's being made on the right basis is expressed counterfactually, the resulting requirements imply that a person will be treated unjustly if her case is decided by a statistical generalization alone, even if the correlation it expresses is perfect in the actual world, that is, even if, in the actual world, 100% of, say, people with engineering degrees are men. This is because such a procedure

⁵⁷ If one takes credences to be dispositions to act then decisions motivated by identity will seem to be evidence of credences based on identity. Still, the decision to offer a job or to listen is not itself a credence, and does not automatically follow from one.

would deny her credibility even if she were qualified, so the decision would not be made on the basis of reliability. Epistemic injustice is that class of violations of the *Fairness* condition and the first *Just Procedure* condition in which the outcome for the speaker depends on a protected identity characteristic.

Justice requires that frequency statistics connecting identities with qualifications not be the basis of employment or listening decisions, but this is not a breach of our epistemic obligation to take all of our evidence into account. The Requirement of Total Evidence can be justified by an assumption that the axioms of probability are rationality constraints on our credences (and seems impossible to justify any other way), but frequency statistics are not credences. Thus the Requirement of Total Evidence does not obligate us to take frequencies into account. Taking frequencies into account in our credences requires direct inference, which, on everyone's view, has further requirements, and when those requirements are fulfilled the inference is generally taken to be permissible and not obligatory. The strict condition imposed on direct inference by Levi is not satisfied in most ordinary contexts because we do not meet people randomly, and they do not circulate randomly. Thus, though Levi justified the condition epistemically, it happens to put ethically salutary limits on our right to jump to conclusions about individuals merely on the basis of facts about populations.

However, Levi's condition permits us to come to a credence .96 that a woman whose folder is randomly plucked from the set of women applicants for the hospital jobs is not an orthopedic surgeon. I have argued that the cause of justice has no need to strengthen the epistemic requirement on direct inference or to classify this belief state as immoral, firstly because a degree of belief and an action are two different things, and the former can justify the latter only when combined with values. The credence about the woman is not the same as and does not by itself justify the action of rejecting her application for the job (by sending her file to the pediatrics department), so to refrain from rejecting the file in that situation, as justice requires, does not count as ignoring evidence. The role of justice is in making a moral demand that you wait for more evidence before deciding whether to remove a woman from the competition because a decision based on frequency statistics would violate *Just Procedure*. But the expected utility paradigm would also counsel waiting as long as you consider the cost in injustice and possible failure to hire the best qualified candidate to outweigh the time and labor spent reading files and the opportunity cost of waiting. If we value just decision-making, then the Principle of Total Evidence will promote it.

References

Adler, Jonathan, "Epistemological Problems of Testimony", The Stanford Encyclopedia of Philosophy (Summer 2015 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/sum2015/entries/testimony-episprob/>

- Alvarez, Maria, "Reasons for Action: Justification, Motivation, Explanation", The Stanford Encyclopedia of Philosophy (Winter 2017 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/win2017/entries/reasons-just-vs-expl/>>.
- AAMC (American Academy of Medical Colleges). *Physician Specialty Data Report 2018*.
- AAMC (American Academy of Medical Colleges). *Physician Specialty Data Report 2012*.
- Ayer, A.J. (1957), "The Conception of Probability as a Logical Relation", in *Observation and Interpretation: A Symposium of Philosophers and Physicists*, ed. S. Koerner, 12-20. London: Butterworth Scientific Publications.
- Basu, Rima (2018), "The Wrongs of Racist Beliefs", *Philosophical Studies*. Published online <https://doi.org/10.1007/s11098-018-1137-0>
- (2019a), "What we epistemically owe to each other", *Philosophical Studies* 176: 915-931.
- (2019b), "Radical moral encroachment: The moral stakes of racist beliefs", *Philosophical Issues* 29: 9-23.
- Brotherton SE, Etzel SI (2019) Graduate Medical Education, 2018-2019. *JAMA*. 322(10):996–1016. doi:10.1001/jama.2019.10155
- Burge, Tyler (1992), "Content Preservation", *Philosophical Review* 102 (4): 457– 88.
- Carnap, Rudolf (1947), "On the application of inductive logic", *Philosophy and Phenomenological Research* (8): 133-148.
- (1950). *Logical Foundations of Probability*. Chicago: University of Chicago Press.
- Carnevale, Anthony P., Jeff Strohl, and Michelle Melton (2011). *What's It Worth? The Economic Value of College Majors*. Center on Education and the Workforce. Georgetown University.
- Cassens Weiss, Debrah (2018), "Lawyer population 15% higher than 10 years ago, new ABA data shows", *American Bar Association Journal*.
- Cherf, John (2019), "A Snapshot of U.S. Orthopaedic Surgeons: Results from the 2018 OPUS Survey", *AAOS NOW*.
- Civil Rights Act of 1991. Pub. L. No. 102-166, 105 Stat. 1071.
- Coady, C.A.J. (1992). *Testimony: A Philosophical Study*. Oxford: Oxford University Press.
- Eisenstadt, Leora F. and Boles, Jeffrey R., Intent and Liability in Employment Discrimination (March 29, 2016). *American Business Law Journal*; Fox School of Business Research Paper No. 16-013. Available at SSRN: <https://ssrn.com/abstract=2756109>
- Fricker, Elizabeth (1995). "Against Gullibility", in B.K. Matilal and A. Chakrabarti (eds.) *Knowing from Words*. Boston: Kluwer., 125–161.

- Fricker, Miranda (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press.
- Friedman, Joel Wm. (2017). *Employment Discrimination*, Third Edition. New York: Wolters Kluwer.
- Gardiner, Georgina (2018), "Evidentialism and Moral Encroachment", in K. McCain (ed.) *Believing in Accordance with the Evidence*. Synthese Library 398, https://doi.org/10.1007/978-3-319-95993-1_11
- Gendler, Tamar Szabò (2011), "On the Epistemic Costs of Implicit Bias", *Philosophical Studies* (156): 33–63. DOI 10.1007/s11098-011-9801-7
- Gertz, Michael, "NFL Census 2016" in ProFootballLogic, <http://www.profootballlogic.com/articles/nfl-census-2016/>
- Good, I.J. (1967), "The Principle of Total Evidence", *The British Journal for the Philosophy of Science* 17(4): 319-321.
- Hàjek, Alan (2007), "The Reference Class Problem is Your Problem Too", *Synthese* 156: 563-585.
- (2008), "Arguments for – or against? – Probabilism", *British Journal for the Philosophy of Science* (59): 793–819.
- Hazlett, Allan (2020), "Critical Injustice", *Pacific Philosophical Quarterly*. Forthcoming.
- Hempel, C.G. (1960), "Inductive Inconsistencies", *Synthese* 12(4): 439-469.
- Joyce, James (2009), "Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief," in F. Hüber and C. Schmidt-Petri (eds), *Degrees of Belief*. Springer.
- Kelly, Daniel and Erica Roedder (2008), "Racial Cognition and the Ethics of Implicit Bias", *Philosophy Compass* 3 (3): 522-540.
- Kidd, Ian James, José Medina, and Gaile Pohlhaus, Jr. (2017). *The Routledge Handbook of Epistemic Injustice*. New York, Oxon: Routledge.
- Korcz, Keith Allen, "The Epistemic Basing Relation", *The Stanford Encyclopedia of Philosophy* (Fall 2019 Edition), Edward N. Zalta (ed.), URL = [<https://plato.stanford.edu/archives/fall2019/entries/basing-epistemic/>](https://plato.stanford.edu/archives/fall2019/entries/basing-epistemic/).
- Kyburg, Henry (1977) "Randomness and the Right Reference Class", *The Journal of Philosophy* 74 (9): 501-521.
- (1983), "The Reference Class", *Philosophy of Science* 50(3): 374-397.
- Lee, Harper (2010). *To Kill a Mockingbird*. London: Arrow Books. (First published 1960.)
- Levi, I. (1963), "Corroboration and Rules of Acceptance", *The British Journal for the Philosophy of Science* 13 (52): 307-313.

- (1977), "Direct Inference", *The Journal of Philosophy* 74 (1): 5-29.
- Levin, Eve A. (2018), "Gender-Normed Physical-Ability Tests under Title VII", *Columbia Law Review* 118 (2): 567-603.
- Macy v. Holder, No. 0120120821, 2012 (EEOC Apr 20 2012), 298.
- Maher, Patrick (1993). *Betting on Theories*. Cambridge: Cambridge University Press.
- McGrew, Timothy (2001), "Direct Inference and the Problem of Induction", *The Monist* 84 (2): 153-178.
- Minghella, Anthony (1999), *The Talented Mr. Ripley*. Feature film based on Patricia Highsmith's 1955 novel (New York City: Norton, W&W & Company).
- Moss, Sarah (2018). *Probabilistic Knowledge*. Oxford: Oxford University Press.
- (2019), "Moral Encroachment", *Proceedings of the Aristotelian Society*, Vol. cxviii, Part 2, IX.
- Price Waterhouse v. Hopkins, 490 U.S. 228 (1989).
- Pritchard, Duncan (2018), "Legal risk, legal evidence and the arithmetic of criminal justice", *Jurisprudence* 9 (1), 108-119, DOI: 10.1080/20403313.2017.1352323
- Reichenbach, Hans (1949). *Theory of Probability*, 2nd edition. Berkeley and Los Angeles: University of California Press. (First published 1934)
- Reid, Thomas (1764). *Inquiry into the Human Mind*, ch. 6, sect. xxiv: "Of the Analogy between Perception and the Credit We Give to Human Testimony."
- Rohde, R.S., J.W. Moriatis, and J.E. Adams (2016), "Where are the women in orthopedic surgery?" *Clinical Orthopedics and Related Research* 474 (9): 1950–1956.
- Smith, Michael (2018), "When does evidence suffice for conviction?", *Mind* 127 (508): 1193-1218.
- Starr, William, "Counterfactuals", *The Stanford Encyclopedia of Philosophy* (Fall 2019 Edition), Edward N. Zalta (ed.), URL = [<https://plato.stanford.edu/archives/fall2019/entries/counterfactuals/>](https://plato.stanford.edu/archives/fall2019/entries/counterfactuals/).
- Sullivan, Charles A. (2009), "The Phoenix from the Ash: Proving Discrimination by Comparators", *Alabama Law Review* 60 (2): 191-239.
- (2011), "Plausibly Pleading Employment Discrimination", *William and Mary Law Review* 52: 1613-1678.
- Suppes, Patrick (1966), "Probabilistic Inference and the Concept of Total Evidence", in J. Hintikka and P. Suppes eds. *Aspects of inductive Logic*. Amsterdam: North-Holland, 49-65.

Thomson, Judith Jarvis (1986), "Liability and Individualized Evidence", *Law and Contemporary Problems* 49: 199-219.

Title VII of the Civil Rights Act of 1964 (Pub. L. 88-352, 42 USC § 2000e [703])

US Bureau of Labor Statistics (2018), "Labor force characteristics by race and ethnicity, 2017", *Current Population Survey 2017*. <https://www.bls.gov/opub/reports/race-and-ethnicity/2017/pdf/home.pdf>

White, Roger (2010), "Evidential Symmetry and Mushy Credence", In T. Gendler and J. Hawthorne (eds.) *Oxford Studies in Epistemology* Vol. 3. Oxford: Oxford University Press, 163-186.

Williamson, Timothy (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.

Wright, Richard W. (2011), "Proving Causation: Probability vs. Belief", *Perspectives on Causation*, Ch. 10, R. Goldberg, ed., Hart Publishing.

Zynda, Lyle (2000), "Representation Theorems and Realism About Degrees of Belief", *Philosophy of Science* 67: 45-69.