

# Bias and interpersonal skepticism

Robert Pasnau

University of Colorado Boulder

**Correspondence**

Robert Pasnau, University of Colorado Boulder.

Email: [pasnau@colorado.edu](mailto:pasnau@colorado.edu)

## Abstract

Recent philosophy has paid considerable attention to the way our biases are liable to encroach upon our cognitive lives, diminishing our capacity to know and unjustly denigrating the knowledge of others. The extent of the bias, and the range of domains to which it applies, has struck some as so great as to license talk of a new form of skepticism. I argue that these depressing consequences are real and, in some ways, even more intractable than has previously been recognized. For the difficulties we face in this domain are fueled not only by implicit biases but by various other sorts of entrenched cognitive attitudes we bear toward others, whether or not we judge them to be our peers. Inasmuch as the epistemic standing of this broader set of attitudes is itself quite dubious, the problem of epistemic injustice turns out to be just one special case—albeit of a particularly nasty kind—from a broader domain of cases where the collaborative character of knowledge clashes with tendencies that make collaboration difficult. This makes the threat of skepticism all the greater, and at the same time makes it harder to see what path of escape there might be.

## KEYWORDS

disagreement, epistemic injustice, implicit bias, self-trust, skepticism

It is a familiar and melancholy fact of life that people in possession of the same information often come to wildly disparate conclusions about matters of the utmost importance. We watch video of the same incidents, we listen to the same testimony, and yet we arrive at opposite conclusions,

often held with passionate intensity, regarding who is at fault and who is telling the truth. Our inability to reach anything remotely like a consensus in such cases might be regarded as the central problem of social epistemology. Yet philosophical attention to these issues has been highly fragmented, with some focused solely on the bare fact of disagreement, and others on the various social biases that distort our epistemic judgment. We have been working our way into this murky terrain from various directions, not realizing that we are all attempting to drain the same swamp.

My own path into the swamp of interpersonal disagreement will run through recent work on epistemic injustice and implicit bias. I will argue, first, that these phenomena raise epistemic concerns that are more vexing than has been generally understood. Then, having identified the properly epistemic structure of these issues, I will argue that they are part of a broader context, the larger swamp. That is to say, the epistemically pernicious influence of racism and other social biases is just a special case from a broader family of interpersonal cases where the collaborative character of knowledge clashes with tendencies that make collaboration difficult.

I begin by arguing that our social biases plausibly give rise to a pervasive skepticism when it comes to our judgments about others. The argument does not turn on excessively credulous worries about implicit bias, nor on tendentially selected marginal cases. On the contrary, the interpersonal cognitive difficulties we face in such cases generate persistent dilemmas of belief from which there is no clear escape. At first glance, such claims may strike some as an overzealous application of identity politics to epistemology. But reflection on these doxastic dilemmas turns out to illuminate just why this sociopolitical territory is so fraught: why facts about social identity call out for recognition and response, even while we often find it difficult to know how best to respond.

Once the structure of the skeptical challenge that arises from social bias becomes clear, it becomes possible to make the generalizing move. The dilemmas we face when it comes to racism and other social categories, far from being unique, are of fundamentally the same kind as those that afflict interpersonal epistemic relationships more broadly. In particular, just as we must evaluate others through various sorts of social biases, so we must evaluate *ourselves*, in light of our own self-biases, in relation to others. Our epistemic lives are not solipsistic; we must constantly decide how to update our own beliefs in light of information from experts, from our peers, and from the remembered judgments of our past selves. Once we grasp these issues in their full generality, a unified field within social epistemology comes into view, unified around the observation that our lives require epistemic engagement with others, and yet we view others through the filter of biases we do not know how to correct for. Whether or not one shares my skepticism over the prospects for our successfully calibrating against these biases, reflection on these issues in their fully generality helps illuminate the broader contours of our epistemic predicament.

In outline: section one introduces the background framework of epistemic injustice and implicit bias, and offers a delimited formulation of the view that such biases have skeptical implications. Section two argues that these implications, even when sharply delimited, are specifically epistemic, yielding a series of dilemmas regarding what we ought to believe. Section three shows how the impacts of bias are prone to multiply over time, which in turn means that the threat of skepticism goes beyond borderline cases. Section four shows how the epistemic issues that arise for social bias can be generalized to include cases of self-evaluation, thus encompassing the seemingly quite different cases of peer disagreement and self-trust.

## 1 | EPISTEMIC INJUSTICE AND BIAS SKEPTICISM

The notion of epistemic injustice needs little introduction, ever since its influential exposition in Miranda Fricker's 2007 book by that name. To be brief, her principal interest is in *testimonial injustice*, where "a prejudice on the hearer's part causes him to give the speaker less credibility than he would otherwise have given."<sup>1</sup> Fricker's particular focus is on injustices that are persistent and systematic, which is to say that they occur frequently across a wide range of social contexts. Accordingly, the principal triggers for epistemic injustice are race, gender, disability, and class.<sup>2</sup> Although Fricker speaks in terms of testimony, subsequent discussions have made it clear that this is best understood quite broadly, so as to encompass any sort of judgments made about the intellectual merits of a person's abilities, actions, or assertions. Even prejudice has come to seem non-essential to the story, inasmuch as epistemic injustice may in some cases be the product of largescale structural explanations, rather than individual prejudices (Anderson, 2012; Haslanger, 2015). Accordingly, Fricker (2013) allows that epistemic injustice is an "umbrella concept" under which prejudice is just one kind of case.

The last decade has seen a rush of careful work on the sort of prejudicial credibility deficit that Fricker describes. The concrete social implications of such work might lead some to pigeon-hole these discussions as a matter of ephemeral politics, marginal to philosophy. Here I aim to show how this judgment is philosophically mistaken. It is worth briefly mentioning, however, that there is also an historical mistake here. To be sure, one is not likely to find epistemic injustice discussed among those authors who enjoy pride of place in the historical canon. But it takes only the most casual perusal of authors on the margins of that canon to find such concerns expressed with poignant force. Here, for instance, is Marie de Gournay, back in 1595:

Blessed indeed are you, Reader, who can be wise without it being a crime, since your sex accords you the privilege of every proper action and speech, as well as the credit to have what you say believed, or at least listened to. As for me, if I wish to put my auditors to the sort of examination that involves, it is said, strings that female fingers cannot touch, even had I the arguments of Carneades, there is no one so much a weakling that he will not rebuke me, to the grave approbation of the company present, with a smile, a nod, or some jest, which will have the effect of saying, "It's a woman speaking."<sup>3</sup>

This describes the basic phenomenon as vividly as one could wish. And if it has taken until the last few decades for academic philosophy at large to take such concerns as worthy of sustained theoretical attention, that, ironically, is itself just a manifestation of the very phenomenon at issue. But as philosophy moves haltingly toward a more inclusive era, crossing lines of gender, race, and culture, the epistemic impact of prejudice has quite properly moved toward the center of our philosophical concerns.

As of now, however, in the short time these questions have flourished, there has been a marked imbalance of attention on the various aspects of the problem. The scholarly literature has had much to say about the ethical and political questions that surround prejudice, and much to say about the psychology of prejudice and bias. From an epistemic point of view, however, there has been relatively little work done, even though the phenomenon is, in its essence, a problem of epistemology.<sup>4</sup> Here I wish to develop the epistemic side of these issues.

One prominent attempt to explore the epistemic consequences of social bias is Jennifer Saul's 2013 paper, "Scepticism and Implicit Bias."<sup>5</sup> Saul builds on Fricker's work, but pays particular

attention to the pervasiveness of prejudice in our lives. Drawing on psychological research on implicit bias, she concludes that we have good reason to worry about the reliability of our judgments “whenever we consider a claim, an argument, a suggestion, a question, etc. from a person whose apparent social group we’re in a position to recognize” (250). Given how much of our rational lives turns on judging the claims of others, and given that we almost always identify people as part of one or another social group, this means that the threat to our reliability is constant. As Saul concludes, “what the literature on implicit bias shows us is that we *really should not* trust ourselves as inquirers” (253). Moreover, Saul argues, these biases are quite ineradicable, at least in the world as it currently exists. To be sure, there are things we can do to lessen our biases,<sup>6</sup> and steps we can take, such as anonymization, to keep our biases at bay. But there is nothing we can do to eliminate such biases wholly, and no basis for confidence that any steps we have taken will be even close to adequate. Her paper’s pessimistic conclusion is that such skepticism will be ineliminable from our lives until we manage “a sweeping and radical transformation of our social world” (261), erasing the conditions that have fostered our biases. All of this leaves us with a form of skepticism—I will call it *bias skepticism*—that is both pervasive and, unlike the usual philosophical skepticisms, based on premises we have good reason to accept as true.

Bias skepticism offers a useful route into the properly epistemic character of epistemic injustice. First, as I will now show, even readers unpersuaded by Saul’s own statement of the case should take seriously a more modest form of the argument. Second, as subsequent sections will show, reflection on the dynamics of epistemic bias over time illuminates the complexity of the underlying epistemic issues. These results have significance even for someone who would reject a modest form of bias skepticism, in part because of their intrinsic interest, and in part because they point the way, ultimately, to a generalized picture of the relationship between bias and interpersonal epistemic relations.

First, then, a more modest skepticism. I myself find Saul’s claims persuasive as written. But readers of a less skeptical cast of mind might consider an argument delimited in three respects. First, we can weaken its key premise: that implicit bias runs rampant through the lives of all of us; that we are all at high risk of error “whenever we are dealing with the social world in a non-anonymised manner” (251). A large body of work, using varied methodology, bears out this conclusion. Even so, some of the details have been the subject of heated dispute, and there are reasonable doubts in some quarters about whether the phenomenon may have been overstated. The state of the literature, then, most of which lies outside the expertise of philosophy, does not presently allow complete confidence in the scope or extent of this bias. Yet Saul’s argument does not require nearly that much. To yield a skeptical result, the argument requires only a well-founded fear that such biases are widespread. Even those who are most suspicious about the implicit bias literature ought to accept, minimally, that the state of research at present raises enough concerns to defeat many claims to justification regarding our judgments of others.<sup>7</sup>

If we cannot trust our judgments regarding the epistemic merits of others, then this will have ramifications across our cognitive lives. Here a second restriction to the argument is in order. Saul makes it clear, as we have seen, that the threat applies to all of our non-anonymized interpersonal interactions. But it is unclear how far these consequences spread. On a narrow reading of her argument, bias skepticism extends only to our immediate judgments about other people. On a wide reading, the skepticism would be far more extensive, encompassing any belief that is the product of some socially mediated interpersonal judgment. Understood in this latter way, bias skepticism might amount to a nearly global skepticism touching on everything we believe in science, politics, philosophy, and so on. For it is not clear that any domain of knowledge would be free from the taint of interpersonal bias, at some point along its path of development. Taken to the extreme, the

only spheres of knowledge left untouched would be what we can know purely through our own resources, unmediated by the contributions of others.<sup>8</sup> To be sure, the narrow conclusion might be the wedge that drives us toward the wide conclusion, since if each of our immediate judgments regarding others is in doubt, then that would call into question the entire web of knowledge. To begin, however, I will focus on the narrow claim, and gradually expand the argument's scope as we proceed.

Once this second restriction is in place, we can make a third restriction, regarding the *degree* to which these direct judgments are unreliable. Here it is important to register a critical difference between bias skepticism and more familiar forms of skepticism. Philosophical skepticisms are ordinarily radical: they question not whether we correctly gauge the size and color of external bodies, but whether there *are* external bodies at all. They doubt not whether we accurately estimate the significance of past events, but whether the world *has* a past that extends beyond the last five minutes. Bias skepticism, in contrast, poses an incremental challenge: not whether there *are* other people, but whether we correctly assess their epistemic merits. Fundamentally, it is a worry about how well calibrated our epistemic judgments are. On its face, this might not seem to count as a skeptical worry at all. After all, it is a familiar fact, in many domains, that our cognitive faculties are far from perfectly calibrated. There is, for instance, well-documented inter-subjective variation in color perception and, indeed, in perception of all sorts. Moreover, it is now widely understood that our cognitive mechanisms rely on a wide array of biases and heuristics that deliver fast and reasonably accurate results. So if it turns out that our biases keep us from being ideally calibrated with respect to our impressions of others, well, this might seem hardly to count as a form of skepticism.<sup>9</sup>

But that is too hasty. There are few if any domains where we expect perfect precision, and so the discovery of small incremental errors is often perfectly tolerable. Extended out to a certain decimal point, there are things we cannot know, and yet no one would speak in such cases of skepticism. Still, if the variation is large enough, and comes in domains where the differences matter, then such concerns may well count as skeptical. If, for instance, we are trying to land a rocket on Mars, and the concern arises that a computer calculation is even very slightly distorted, then that is obviously a difference that matters. Similarly, if we think that an admissions committee is implicitly biased toward certain sorts of candidates, then even if that bias is slight, it may be enough to influence the ultimate rankings, and so determine who is and is not admitted. In contexts such as this, skepticism seems warranted.

Although bias skepticism gains plausibility through these three restrictions, the price is a serious diminishment of its scope. By restricting our attention to *incremental* discrepancies in our *immediate* judgments about the worth of others, we can no longer lay claim to the sorts of nearly global impacts that Saul describes, as when she remarks that “implicit bias could be affecting one’s reasoning at almost any point—it is very hard to judge when social group membership is having a pernicious influence” (258). Delimited bias skepticism does not go nearly so far. It makes no claims about the remote ramifications of bias in one’s reasoning, and it does not suggest that incremental differences have global epistemic impacts. Instead, my immediate focus will be on one specific family of cases, those involving *contested interpersonal judgments*, where two or more people are subject to competitive evaluation. Inasmuch as this is not something we do all the time, the skepticism at issue is localized. Still, even in this delimited form, bias skepticism casts a large shadow over our lives. We constantly make interpersonal judgments about others, just as we are constantly judged. Success in life turns on such judgments and, accordingly, for those in a position of authority, the correctness of these judgments should be a matter of the utmost concern. In what follows, I do not ask that readers be fully ready to *endorse* bias skepticism, even in this modest

form. All I require from this first stage is a reader who is prepared to treat bias skepticism as a serious threat. Indeed, my ideal reader at this point is not someone who has wholly surrendered to skepticism, but someone who, like the Pyrrhonian skeptics of old, continues to inquire hopefully into the prospects for a solution. That process of inquiry is what will lead to a broader picture of contested interpersonal judgments and the difficulties that surround them.

## 2 | DOXASTIC DILEMMAS

The modest form of bias skepticism I have sketched retreats from bold empirical claims about the psychological impact of bias, and acknowledges that these impacts may be incremental and narrowly focused on contested interpersonal judgments. It is in these cases that the subject's properly epistemic character is most salient. After all, epistemic injustice that is the product of *explicit* prejudice is so egregious morally as to overshadow whatever epistemic failings may be involved. The racist juries and sexist coworkers that populate the most familiar examples require, first and foremost, moral enlightenment; their epistemic deficits look derivative next to their obvious moral failings. Shifting to subtler forms of bias both provides the fuel for Saul's form of skepticism and puts the properly epistemic aspects of epistemic injustice on center stage.

Yet it may seem at this point that bias skepticism has been so delimited as no longer to be a serious problem within epistemology. First, in its localized form, it may seem to present only familiar practical worries about procedural fairness rather than distinctly epistemic worries. Second, to the extent the worries are epistemic, they may seem to arise only in the sort of borderline cases where a healthy dose of skepticism looks quite unsurprising. Taking these concerns in turn, I will argue in this section that the threat of bias skepticism, even in its delimited form, gives rise to two specifically epistemic dilemmas over what to believe—hereafter, *doxastic dilemmas*. The following section will then argue that these dilemmas go beyond borderline cases.

The first of my two doxastic dilemma arises over whether we ought to *suspend* beliefs of the relevant sort; the second concerns whether we ought to *recalibrate* them. To begin with the first, it is an old and familiar thought, running all the way back to antiquity, that skepticism gives rise to a practical dilemma over the implications for belief. From one side, it seems that where we do not know, we should suspend belief. From the other side is the equally obvious thought that suspending belief, in any thoroughgoing way, would leave us incapable of navigating the world as is required in a meaningful life. Does this old dilemma arise out of bias skepticism? Suppose that you are serving on an admissions committee, and you must rank the top 20 candidates from among a pool of 200. Certainly, there is no pressure to withhold belief about whether the strongest-looking candidate is better than the weakest-looking candidate; nothing requires us to suppose that our judgments are *that* badly calibrated. At a certain point, however, the skeptical argument will begin to have real force, in view of the fine-grained evaluative judgments that must be made. Still, it would be reasonable to argue there is no dilemma here, because the properly skeptical committee member need not make any commitment to the objective *correctness* of her choices. She is doing the best she can, and she hopes that her colleagues are all doing the best that they can, and so she trusts in the process to yield a ranking. Does she believe the ranking is objectively correct? Not only is there no need for her to do so, but, moreover, there would be a kind of zealotry in the opposite direction if she were actually to *believe* that the committee managed to produce the uniquely correct ranking of candidates. In a case like this, belief can be replaced by what, borrowing a label from political theory, we might call *pure procedural justification*. This is not the

sort of justification that warrants belief, but why should one form beliefs in such a case? The old skeptical dilemma seems, at first glance, to gain little traction here.

Given the localized character of delimited bias skepticism, abandoning belief may seem like the obvious way forward. But militating against a general suspension of belief is the way these sorts of interpersonal assessments often occur in contexts in which we need to enter into close interpersonal relationships with those we assess. Although, in cases such as this, bias skepticism gives us reason to doubt our judgments about the worth of others, such doubts are not always consistent with the demands of friendship and collegiality. As the recent literature on interpersonal relationships has stressed, such bonds demand a certain willingness to *believe in* other people, and specifically to believe in their talents in the relevant fields where their worth is contested.<sup>10</sup>

To see how such pressure to believe can build in the cases before us, consider more fully the context of graduate admissions. A professor who does not believe in the accuracy of the admissions process might, without any harm done, tell a disappointed student not admitted off the waitlist that the process is highly fallible and somewhat random. And it might be positively healthy to say that same thing to the happy students who have been admitted. But now let a year go by and imagine that formerly happy student in your office, unhappy about your negative comments on his paper. Here too bias skepticism tells us that we cannot have firm confidence in our evaluations of others. But now it seems far less acceptable to tell the student that the grading process—just like the admissions process (!)—is highly fallible and perhaps even a bit random. Here it does not seem enough to rely on pure procedural justification: that you evaluated the student as best you could, but that you admittedly have only a certain level of confidence in your judgment. A sustainable relationship between teacher and student requires more than acting out that relationship in accord with the correct rules of a procedure. Even if teachers are capable of functioning while abstaining from judgments about the objective merits of their students, the result is hardly likely to be the sort of engaged and fostering mentorship that students deserve.

Such tensions arise only in certain kinds of cases. Where I am under no obligation to make a fine-grained determination of merit, or where there is no close personal relationship involved, I will feel neither skeptical pressure nor pressure to form a belief. For instance, musical amateur that I am, I can just happily sit in the audience and enjoy the second violinist's playing. Matters are very different, however, for those who must sustain a close and discriminating interpersonal relationship. The other members of the quartet, let us imagine, had to choose a violinist from among fifty others candidates. For the duration of the search, this required certain procedural commitments, but extending beyond that is the pressure to enter into a doxastic commitment: to believe in their new second violinist. Even given the fine-grained discrimination required among candidates and the ever-present risk of bias, there would be something problematic about the members of the quartet not fully believing in the correctness of their choice.

We have seen, then, how the interpersonal character of bias skepticism imposes pressure from one direction, making the suspension of belief in some cases problematic. This all by itself creates something of a dilemma, by pushing back against the familiar skeptical pressure toward suspension of belief. But matters are still more complex in the context of bias skepticism, because the judgments at issue have a significant moral dimension. When we underestimate a person's worth because of race, gender, disability, or class, we are doing something that is seriously morally wrong. Suppose the violinist who gets the job is a white man. Was he the best candidate, or did he get the job because of biases in the selection process? Procedurally, one does the best one can, and lives with the result. But should one also *believe* these results? That seems like an optional further commitment, which one can avoid by embracing bias skepticism and accordingly giving up

one's beliefs in such cases. The usual epistemic pressure to abandon belief in the face of skeptical threats ought to have all the more weight in cases where the moral stakes are high.

This moral pressure might come in either of two ways. On the most straightforward account, moral concerns give us non-epistemic reasons for belief. Beliefs may be wrong, on this line of thought, because they are morally wrong—e.g., because they are racist—quite apart from whatever their epistemic status may be (Basu, 2019; Gendler, 2011; Pasnau, 2018). So understood, the moral cost of biased judgments would give us an added reason to suspend belief, beyond the more familiar epistemic costs associated with the failure to track what is true. Inasmuch as the belief dilemma need not be understood as purely epistemic, this additional moral cost may count for just as much as the epistemic cost.

Instead, or in addition, one might think that moral considerations encroach directly on the epistemic domain. Just as it is often thought that heightened practical stakes make a difference to what we can know, and thereby make a difference to what we ought to believe, so it has recently been argued that heightened moral stakes introduce higher standards for knowledge and thereby for belief (Basu & Schroeder, 2019; Pace, 2011). This line of thought locates the moral cost of belief in a different place, among the properly epistemic considerations. Indeed, it gives us a powerful further reason to take bias skepticism seriously. For the encroachment of morality into epistemology would, in effect, have a multiplier effect on the argument from implicit bias, making it even harder to see how we can have knowledge when we judge other people in a socially embedded context.

At this point there will be readers who think the way out of such dilemmas is to abandon the notion of full belief in favor of graded beliefs, or credences. Instead of posing an all-or-nothing question of whether or not to believe, we should simply proportion the degree of our belief to the strength of the evidence. And this is liable to seem especially appealing in the sorts of incremental cases at issue. But this familiar strategy faces special difficulties in the case of bias skepticism. From one side, the features of interpersonal relationships that put pressure on us to believe can scarcely be addressed by carefully chosen credences. Thriving interpersonal relationships require a full and mutual belief in the other's worth, not degrees of partial belief. From the other side, the morally loaded character of these judgments is not the sort of evidential consideration that ought to shift our credences. Rather, such considerations go directly to the question of whether or not to believe, leaving credences untouched.<sup>11</sup> Here too, then, bias skepticism, even in its delimited form, gives rise to a belief dilemma that is significantly more intractable than in ordinary cases of skepticism.

Let us now turn to the second dilemma. It may have seemed perverse to suggest that the moral seriousness of prejudice puts pressure on us to suspend belief in doubtful cases. Instead, it might be thought, the obvious way forward is to give the benefit of the doubt to those who are most at risk. If bias skepticism puts us in a position where we can neither believe our judgments about others nor suspend belief, then perhaps we should give ourselves some reason for confidence, by *recalibrating* our judgments. Admissions committees might thus give extra weight to students from certain social groups, and less weight to students from other groups. Teachers might similarly recalibrate their grades.

To ensure that our concerns remain epistemic, let us again register, only to set aside, the level of pure procedural justification. From a practical point of view, we can grant it will sometimes be appropriate to give extra weight in these ways. Our question is whether recalibration offers a response to the belief dilemma, by allowing us to believe our judgments in these contested interpersonal cases while still taking our biases seriously. This sort of compromise solution to the first dilemma, I now want to argue, creates its own dilemma, over the very question of whether we





FIGURE 1 The Ebbinghaus-Titchener illusion [Color figure can be viewed at [wileyonlinelibrary.com](http://wileyonlinelibrary.com)]

ought to recalibrate. Consider this concrete example of recalibration in action, from Daniel Kelly and Erica Roedder (2008, 533):

[S]uppose you are a White professor grading a Black student's paper, and you are initially inclined to give the paper an 89/100. Does the possibility of implicit racial bias give you good reason to think the paper actually deserves slightly better, e.g., 90 or 91 points? .... An analogy will be helpful here. Suppose you learn of psychological research showing that most people are inclined to underestimate the size of circles when set across a hatched background. Suppose you are later asked to judge the size of a circle on a hatched background. In deciding the size of the circle, it is most rational to estimate it to be slightly larger than you are initially inclined to guess.... We maintain that by parity of reasoning, it would be wise to make a similar adjustment for the implicit bias in grading, just as you would correct for the visual bias in judging the size of a circle. In both cases, one is acting for purely *epistemic* reasons; in order to give the most accurate grade, i.e., in order to grade the paper based on its merits, it is reasonable for the savvy grader to correct for the effects of racial biases.

Is recalibration the right course of action for Professor White (as we can call him)? Let us accept, as the final sentence stipulates, that his only goal is to grade the papers accurately. This helps make clear that our concern is fundamentally epistemic. Let us further suppose, going beyond what Kelly and Roedder stipulate here, that implicit racial bias is in the imagined case not just a "possibility" but is in fact very likely. Given those assumptions, the case for recalibration is obvious. For even if the paper looks to Professor White like an 89, still he has good reason not to trust his judgment, and to give it a higher score. Analogously, a naive observer of figure 1 is very likely to underestimate the size of the inner circle on the left.

You should recalibrate, and judge the left inner circle to be larger than it looks. On the face of things, Professor White should do the same.

Strengthening Professor White's case for recalibration, at least arguably, is the above phenomenon of moral encroachment. For, given the badness of biased misjudgments against disadvantaged social classes, it is natural to suppose that errors, if there must be any, should be made to favor the disadvantaged. The moral pressures that give bias skepticism special force steadily recede, one might say, the more we recalibrate away from our biases.<sup>12</sup> Yet, though I would not deny that there is a moral asymmetry here, such considerations have their limits. Exaggerated corrections for bias may create their own harms, even for those in whose favor they work (Zheng, 2016, 81–82), and of course recalibrating in one direction will often require an

inverse recalibration in another. To admit one student off the waitlist, for instance, means not admitting another. Restricting ourselves to the question of what to believe, it may be less obvious that we are dealing with a zero-sum game: why not, when in doubt, always evaluate the worth of others more generously? But in many evaluational contexts it is not enough just to have a generalized belief in the worth of someone. If the quartet is to recalibrate their belief about the *best* person they could hire, this entails downgrading another finalist. To do so on the basis of the candidate's race, gender, etc., however well intended, carries special moral risks.

These risks notwithstanding, recalibration looks attractive. Yet, when we bear in mind the challenge of bias skepticism, it is hard to see how recalibration actually helps with the dilemma over what we should believe. In the illusion case, we can accurately measure the two circles, and even if we do not ourselves take the time to do that, the ease of doing so gives us reason to trust others when they tell us that the two circles are equal in size. Obviously, however, nothing like this will settle the problems of bias skepticism. Although we have good reason to think that implicit bias is very likely to play a role in grading, we have no way of knowing exactly, or even very roughly, how much of an effect there is in any particular case. Hence any attempt at recalibration runs the risk of being too much or too little, and will inevitably be too crude to reestablish any confidence in our fine-grained judgments. As Kelly and Roeder tell the story, Professor White might add a point, or maybe two points, thereby *doubling* his estimate of his own bias. And is even this enough? Might it be too much? Even a single point would cross the 90% threshold, and so presumably elevate the paper into the A range, making a one percent recalibration more significant than the raw numbers suggest. Moreover, might Professor White *already* have tacitly recalibrated in the student's favor, along the way, in assigning the original 89?

Just as recalibration fails to defuse the threat of skepticism, so it fails to help with the belief dilemma. For let Professor White recalibrate and give the paper a 91. Or, reverting to an earlier example, let an applicant be boosted up to ninth place rather than eleventh, and so be admitted to a PhD program. As a matter of pure procedural justification, these may be the right decisions, with which we can comfortably live. But can we *believe* these outcomes? After all, given the admittedly crude character of the recalibration, the paper might just as well be a 93 or an 88. The student might more deservedly be ranked fifth or fifteenth. On the available evidence, given the range of possibilities, the paper is much more likely *not* to deserve a 91, and the student is much more likely *not* to be the ninth best in the applicant pool. One ought not to believe  $p$  while judging  $\sim p$  to be more likely. But if all of this makes it seem obvious that we should not believe fine-grained results such as these, then again we come up against the fact that we are dealing not just with abstract propositions, but with real people whom we need to respect and instill with confidence. In cases such as that, failing to believe seems equally problematic. So even if recalibration is defensible at the level of pure procedural justification, there remain distinctively *doxastic* dilemmas.

Suppose, nonetheless, idealizing considerably, that Professor White has strong evidence that this particular recalibration makes for an improvement in accuracy. Even so, it would remain doubtful whether he ought to recalibrate. The trouble comes from the disparity in sources of evidence. In favor of his original verdict is that he read the paper carefully and, relying on his expert training and years of experience, came to a considered judgment about its value. The paper *looks to him* like an 89. He can articulate its strengths and weaknesses, and explain why it looks definitely weaker than another paper that received a 91. We might say, then, that Professor White is entitled to trust his own judgment in such a case. Of course, it is precisely here that the evidence of implicit bias comes into play, countering that alleged entitlement. But this counterevidence provides no reason for recalibration that is specific to the paper in question—all his reasons specific to the paper still point to an 89—and indeed the evidence of bias is not even specific to him as an

individual, let alone to his assessment of student work. Even if we stipulate that the evidence of bias is strong, it is not clear that it should get the same kind of consideration that Professor White gives his immediate impressions about the paper's merits.

We face here a conflict between individual and statistical evidence. A large literature in law and philosophy has explored this distinction, and studied why, in certain contexts, we do not give statistical evidence the weight that it would seem on its face to deserve.<sup>13</sup> In the legal context, for instance, if a defendant is charged with sneaking into a concert without a ticket, eyewitness testimony would count as strong evidence, even if it were judged only 90% reliable. But evidence of a purely statistical nature—for instance, the bare fact that only 10% of the people at the concert had purchased a ticket—would likely not be given the same weight, and perhaps not be allowed at all. There is dispute over exactly why such distinctions are appropriate to draw, and where we should draw them, but the intuitive correctness of the distinction seems clear both in law and in the broader epistemic context. Suppose, switching examples, that Professor White, having read an anonymized student essay, judges that the paper is very likely plagiarized, and suppose his instincts about this have been quite reliable over the years. Compare a non-anonymized case, where he forms his belief not on the basis of the paper's content, but simply from knowing that the student belongs to a social group—say, members of a certain fraternity—in which cheating is quite prevalent. Even if, statistically, that fraternity's track record is precisely as bad as Professor White's plagiarism instincts are good, still it would seem obviously reasonable for Professor White to believe the charge of plagiarism in the first case, and obviously wrong in the second case. Recalibrating for racial bias seems to have essentially the same structure, involving a contrast between individual and statistical evidence.

The cases where statistical information seems especially problematic turn on using it as positive evidence to justify a belief. Where such information is instead used indirectly, to shape the direction of inquiry, it is generally less problematic. It is statistical information that tells us, for instance, to be less than fully confident in eyewitness testimony, and information about bias should similarly lead Professor White to reassess his method of evaluating papers. Instead of formulaically recalibrating, he might seek further evidence by rereading the paper, asking someone else to read it, comparing it directly with a comparable paper, evaluating all the papers on a defined template, and so forth.<sup>14</sup> In principle, the strategy of gathering more evidence provides a path through our two doxastic dilemmas, inasmuch as it offers the prospect, in the end, of a recalibrated belief that satisfies high standards for justification. Indeed, this is the path that bias skepticism itself endorses, inasmuch as the bias skeptic does not counsel abandoning all hope, but urges redoubling our efforts to do better. Of course, bias skepticism also tells us that there is no easy route to this ultimate recalibration, inasmuch as the process of acquiring further evidence will be problematic at every step of the way. Here is where it particularly matters that the form of skepticism at issue has been so delimited. As long as we are considering only incremental differences in a narrow class of cases, it is easy to think that bias skepticism is a fairly marginal phenomenon, manageable through a mix of intellectual humility and methodological care. It is time, then, to consider how these marginal cases are liable to expand and intensify over time.

### 3 | BEYOND MARGINAL CASES

So far, I have been focused on a narrow set of cases where the risk of bias is notably large. Yet, by delimiting the scope of bias skepticism in this way, its impact is obviously curtailed. In particular, it is natural to think that the doxastic dilemmas in question arise only from pressure to make

extremely fine-grained judgments. No one, surely, should expect knowledge about the precise worth of a paper on a 100-point scale, or the exact ranking of a candidate among a pool of 200. Indeed, quite apart from worries about social bias, it may be that no one should form beliefs at that level of precision, even if sometimes we have to act on the basis of such fine-grained distinctions.

In this section I will argue that the precisely quantitative and marginal examples considered so far can, over time, expand to yield dramatic qualitative differences in judgment, making the task of recalibration even more difficult and making the belief dilemma that much more serious. For even if it is easiest to recognize the influence of social bias in fine-grained, incremental cases, in fact such biases can have quite dramatic effects. Although empirical research on the cumulative impact of bias has yielded mixed results with regard to the severity of these impacts, there is a sense in which the issue is beyond dispute. After all, as I remarked at the start, we constantly observe that people's judgments regarding others are riven by biases of one sort or another, and not just in incremental ways. People viewing testimony from one or another social or political perspective are prone to judge that a witness is not just marginally more or less likely to be telling the truth, but instead is *obviously* lying or telling the truth. People watching the same video of a police shooting are prone to come to radically different conclusions about who was at fault, in ways that obviously correlate with their antecedent perspectives. We know that this sort of polarization is common, even if we might like to think that we ourselves are largely immune to bias.<sup>15</sup>

As familiar as cases of this more dramatic kind are, it will be worthwhile to develop a particular case in detail. Following the familiar injunction to write what you know, let me return to Marie de Gournay, and imagine her as a faculty member within a modern university, at a department meeting where a vote will be held on whether to hire a male candidate to a faculty position. Suppose that Marie has raised concerns about the candidate's past conduct with regard to women, and suppose that those suspicions have been rebutted by Tony, who testifies to the candidate's upstanding character. The department must vote.

We will want to consider Marie's own perspective, but first let us consider someone, call him Adam, in the position of having to decide between Marie's and Tony's competing claims. Adam could be just blatantly sexist, and accept Tony's testimony over Marie's on that basis, leaving us with an obvious case of epistemic injustice. But the situation becomes epistemically more interesting if we suppose that Adam is, to the best of his ability, a fair and conscientious colleague. He feels no bias toward either Tony or Marie, and listens carefully to their arguments. Now suppose that Adam finds Tony's story to be much more plausible than Marie's. It might well seem to him just bizarre that anyone (especially anyone with a PhD!) would engage in the sort of behavior the candidate is accused of, whereas he takes himself to have seen how false rumors about such things can spread. This might, of course, be the end of Adam's reflections, but suppose he has read about implicit bias; suppose he has even gone to a training session on the subject. He might then consider how bad it would be, morally speaking, if his views were to be unwittingly biased against Marie. So he might then consider recalibrating his judgment. But should he go so far as to vote against the candidate, simply because of the statistical likelihood that he is biased in Tony's favor? Indeed, should he recalibrate at all? After all, it continues to *seem* to him as if Tony's claims are more plausible, and that a no vote would be an injustice to the candidate. As in the case of Professor White, there are moral risks on both sides of this question.<sup>16</sup>

So far, this is simply a more involved example of the sort of incremental case considered in the previous sections. But Tony's reflections might go farther. For he might weigh not only the risk of present bias, but also the influence of a lifetime of bias. After all, his inclination to find the candidate credible is the product in part of the sorts of experiences he has had over his whole life, and he might wonder whether that leads him both to underestimate the likelihood of sexist behavior

and to be overly confident of his own perspective.<sup>17</sup> Moreover, beyond his personal experiences, Tony has been absorbing information from other people for his entire life, and judging some of those sources to be more credible than others. The incremental biases considered in the previous sections do their work here, and these impacts accumulate, as later judgments take shape based on earlier judgments. Since what looks credible depends in large part on Adam's prior beliefs, the impacts of bias can have exponentially greater force as time goes on. So, suppose it looks to Adam as if, time after time, Marie has been overly quick to complain about one or another perceived injustice, whereas Tony has been appropriately skeptical. Might this perspective of dramatic difference be the product of incremental biases the force of which has multiplied over time?<sup>18</sup> In such circumstances, epistemic injustice can take on a complex temporal dimension, where one's current judgment must be weighed not only against the threat of present bias, but also against the risk that one's current judgments are distorted even more dramatically by a history of bias.

Now Adam should be worried not just about bias at the margins, but about the prospect that his judgment might be quite significantly distorted. At this point the verdict of the previous section—that statistical information about bias should lead us to seek more evidence—looks much less helpful than it previously seemed, because if our judgments are as distorted as Adam fears, then more evidence just may not do any good. And, as I have stressed already, we know that this is no idle philosophical possibility, but the real-world circumstances of many people who find themselves locked into a worldview from which no amount of evidence seems able to dislodge them (Bramson et al., 2017; Singer et al., 2019). Of course, *we ourselves* are not like that, or so we think. Yet, going back to Adam, what should *he* believe? It still seems to him as if the candidate is being unjustly accused. In light of all these uncertainties, might it be reasonable for him simply to trust himself? Everything we know about bias suggests, to the contrary, that we should not trust ourselves. But then who or what should Adam trust? Should he, across the board, simply embrace Marie's perspective on all such matters? That's absurd, if for no other reason than that there will be other female voices, with other points of view, and other forms of bias to be considered. Should he, then, just embrace skepticism, and suspend judgment?

As in other cases, it may be clearer what Adam should *do*. Voting is not believing, and so even if he ultimately finds himself in the grip of hopeless skepticism, he may still be able to vote. Relying on pure procedural justification, he can consider which vote is more appropriate to cast under conditions of uncertainty and, without forming a belief as to where the truth lies, he may decide to trust in the collective wisdom of the faculty to determine a fair outcome. Yet, as before, this leaves the doxastic dilemmas untouched, and now that we have gone beyond marginal cases, the difficulties are all the greater. Perhaps Adam need not reveal how he voted, or why, but he has to function going forward as a member of his department. An attitude of unsettled doubtfulness towards all of one's colleagues is hardly the foundation of a healthy professional life. And even if one can imagine sustaining this outlook in one's professional relationships, the threat of bias skepticism is pervasive enough to extend over every aspect of one's life: students, friends, family. The inability to trust one's own judgments, now not only in marginal cases, but even on fundamental matters, entails an inability to trust others, and so directly threatens one's ability to form meaningful interpersonal relationships. You need not have confidence that your student is in the top 7% of his peers, rather than the top 8%, but if you have doubts over whether he belongs in a PhD program at all, then that creates a problem of a specifically doxastic kind. The cumulative impact of bias, over a lifetime of experiences, can yield just this sort of more dramatic uncertainty.

And what of Marie? Ironically, she herself may come to be in a similar epistemic predicament. We have seen how Adam, on careful reflection, might be reduced from an unreflective self-confidence to a wavering uncertainty over whether he can trust his own judgment. For Marie,

this may have long been the usual state of affairs. One of the likely consequences of epistemic injustice for her, over time, is that she internalizes the diminishment of her own credibility that arises from the prejudice of others (Jones, 2012). So whereas Adam, prior to extensive reflection, might be expected to have an artificially elevated sense of confidence, Marie might conversely be excessively doubtful about her own judgment. We can imagine that she might, at an early stage in her intellectual development, be quite unaware of the lack of self-trust that distorts her thinking, then grow into an awareness of it and even achieve some ability to modulate it effectively. But this capacity for self-control will not solve the problem of knowing how to recalibrate: she needs to be somewhat more self-confident, but there is always the risk of overcorrection, which might serve to undermine her further in the eyes of others. Finally, even if she forms a view about exactly how far to recalibrate, the question may constantly arise as to whether she should recalibrate in the present instance. At the meeting, Tony seems so confident, and she feels not nearly so sure. Others, like Adam, seem to find Tony persuasive. Is this bias on their part? She thinks it is. Should she be more confident? Yes, she thinks she should. But, still, with regard to *this* particular claim she is asserting, she *does* genuinely feel somewhat unsure. She wishes to do an injustice neither to the candidate nor to herself. So should she really adhere to a claim with greater confidence than she feels the claim itself deserves? Like Adam, she may find herself facing the accumulated force of doxastic dilemmas from which there is no easy escape.

Marie's predicament can be extended into many domains. After all, the dilemmas generated by epistemic injustice arise not only for those who *give* the grades and *make* the decisions about hiring and admission, but also for those on the receiving end of these judgments. Whether or not one gets admitted to a program, whether or not one gets an A, the question can always arise, regardless of one's social group, of whether the decision was the product of bias, or the product of an attempt to calibrate against bias. Anyone at risk of being the target of such biases may come to doubt their own worth, regardless of whether they are at risk of being victimized or benefited. When these doubts arise only at the margins, the consequences are perhaps not so severe. But once we see how the impact of marginal cases multiplies over time, it becomes evident how bias skepticism can lead to a crisis of belief for anyone whose life has been subject to the potentially distorting influence of social biases.<sup>19</sup>

#### 4 | INTERPERSONAL BIAS GENERALIZED

We have now seen how the threat of bias skepticism arises, why it has serious practical consequences, and how its effects can intensify over time. In light of this detailed account of the structure of the phenomenon, we are in a position to take the final, generalizing step, and see how epistemic injustice is just one particularly vivid and confounding kind of case from a much larger class. For although the examples so far have been driven by social prejudices, the underlying roots of the problem are much more general. Interpersonal skepticism threatens whenever we find ourselves required to make a judgment about the comparable worth of two individuals in one or another domain: about their honesty as witnesses, their talents as musicians, their potential as scholars, and so forth. Of the various prejudices liable to distort such judgments, racism and sexism are the most toxic examples, but the phenomenon extends beyond the domain of these and other social biases. To understand the epistemic issues in their full generality requires this broader perspective.

The breadth of the issues at stake becomes particularly clear when we look not at cases where one person is assessing the worth of two or more others, but instead at first-person cases, where

one makes a comparative judgment about someone else in comparison to *oneself*. These cases have the same structure as the contested interpersonal cases previously considered, except that now there is an element of self-evaluation. Of course, all the same social biases may play a role here: insofar as I inevitably conceive of myself as having a social identity, I may have biases toward myself, positive or negative, which will influence my judgment. But what makes these first-person cases particularly interestingly is that a whole further class of biases enters into the story, inasmuch as comparative judgments of this sort must take into account self-bias—that is, our natural prejudices in favor of our own perspective.

Consider, to take a familiar example, a case where I disagree with someone whom I acknowledge to be my epistemic peer. It is clear that equally intelligent and well-informed people come to varied conclusions about a great many matters—in politics, religion, philosophy, and so on—and yet most of us retain our own convictions even in the face of this extensive disagreement. We think, presumably, that we have good reasons for those convictions. Yet at the same time we recognize that others, who in this domain we have every reason to judge our equals, reject those reasons and have reasons of their own. Recent philosophy has been sharply divided over whether, in cases such as this, it is rational to remain steadfast in our own views, or else become more conciliatory toward the views of others (Feldman & Warfield, 2010).

Although the existing literature treats peer disagreement and epistemic injustice as wholly distinct topics, we are now in a position to see how intimately related they are. Each characteristically involves a contested judgment of interpersonal worth, where that judgment is threatened by one or another sort of bias. Unsurprisingly, given these structural parallels, the same sorts of considerations that support bias skepticism have led some toward skepticism in cases of interpersonal disagreement.<sup>20</sup> This in turn raises the dilemma of belief: should one follow the evidence of self-bias and suspend judgment, or retain the beliefs that strike us personally to be correct? The dilemma of recalibration arises accordingly. Just as Professor White has to weigh his own impressions of a paper's worth against the threat of racial bias, and decide on that basis whether to recalibrate, so Professor Steadfast, as we might call him, has to weigh his own reliability against his peers, and decide whether his own scholarly judgments ought to be recalibrated in light of the risk of self-bias. Interpersonal considerations again play a role, inasmuch as our beliefs in politics and other matters connect us to friends and family. One's relationship to oneself also enters into the story at this point, inasmuch as the thoroughgoing surrender of one's own perspective might undermine one's very self-identity.

The similarities here are not just structural but substantive, inasmuch as the same epistemic considerations ground the arguments for each approach, steadfast and conciliatory. The standard arguments for steadfastness in the face of peer disagreement are that it is required by a proper degree of self-trust and a proper respect for the first-order evidence that motivates one's own view.<sup>21</sup> These same considerations, as we have seen, motivate steadfastness in cases of social bias. In both cases, moreover, the impacts of bias are liable to multiply over time, making us gradually more entrenched in our biased worldview. In addition, the comparison illuminates something not previously noticed in the disagreement literature: that the evidence in favor of conciliation is purely statistical. Professor Steadfast's reasons for refusing to recalibrate have nothing to do with the content of his opponents' arguments. After all, he takes himself to have good reasons to reject those arguments. The argument for conciliation, instead, is statistical, resting on the idea that, over time, neither of two peers is more likely to be right than the other. Given that peer disagreement is liable to give rise to doxastic dilemmas very much like those generated by social bias, it is no surprise that the literature in this area has arrived at such an impasse.

To be sure, social bias has a moral dimension that is lacking in self-bias. But even here there are instructive commonalities. For one might suppose that general ethical norms of concern for others imply a duty to respect the opinions of others. Again, however, there are countervailing pressures, since one might argue that there is value in self-trust, and disvalue in being too ready to abandon one's own judgments in the face of disagreement. Of course, such concerns are magnified when the lack of trust in others is driven by social biases. Hence there may seem to be more pressure on Professor White to recalibrate his grades in light of racial bias than there is on Professor Steadfast to recalibrate in light of self-bias. But it is perhaps significant that, in the peer disagreement literature, the places where steadfastness seems most appropriate are precisely those cases that are of greatest ethical import: in religion, politics, and philosophy. It is as if the proponents of steadfastness judge these matters to be of too great moral concern to be determined by mere statistical considerations. A similar thought may drive those who think that Professor White ought not to recalibrate: after all, assigning grades in a way that tracks merit is itself a significant moral principle. Similarly, it is a serious moral matter to charge a job candidate with wrongdoing, and one might think this gives Adam strong reasons not to rely on purely statistical information about bias. Yet, once again, there are countervailing considerations. After all, forming a conclusion under the influence of sexism is very bad. In all these cases, if we suppose that moral considerations encroach on the epistemic, then we have to acknowledge that the encroachment comes from all sides. And looming over this whole family of cases is the puzzle of why statistical information gets less weight in certain kinds of cases. For one might suppose that, the greater the stakes in a situation are, practical and moral, the more important it should be to rely on all the available information. I am not claiming that skepticism is the inevitable conclusion, but I would suggest that we are now in a better position to understand why these issues have seemed so intractable.

Peer disagreement is just one special, idealized case of interpersonal disagreement, and so the considerations just surveyed apply equally well to any situation where the special weight given to one's own perspective needs to be weighed against the view of another, regardless of whether the other person is judged to be more, less, or equally reliable. And even here we might generalize still further, because self-trust has a temporal dimension. As a result, beyond comparing our judgments to those of others, we can make *intrapersonal* comparisons, weighing our current confidence against our previous or future confidence. In many contexts, we recognize that our past judgments are likely to be at least as reliable as our current ones. I might, for instance, run across old lecture notes, written up at a time when I was keener and more conscientious. The odds are that those old notes are more reliable than anything I might do today. Yet even when I recognize the odds for what they are, I may still put more trust in how things strike me right now. Our current seemings have a grip on us in a way that our past judgments do not.<sup>22</sup>

Here too, there are close structural and substantive commonalities with cases of epistemic injustice. In light of our known bias in favor of our present selves, we have statistical reasons to recalibrate our present confidence downward. Yet we are reluctant to let considerations of that form override our present judgments of how things seem to be, and considerations of self-trust again push us toward remaining steadfast. Here the role of interpersonal relationships gets wholly replaced by the importance of self-integrity. Accordingly, moral considerations may not arise so obviously in such cases. But, then again, they may. When Amy Winehouse sings, "I cheated myself, like I knew I would," she suggests the way in which this sort of self-bias can amount to an injustice—epistemic and otherwise—done to oneself, over time. A case like this may even feature various other biases from among those we have considered. One may, for instance, neglect one's past self by giving inappropriately large weight to the current views of others. And if one's



social prejudices entail an excessively negative (or positive) self-image, then a case like this may also involve social biases.

There is, moreover, every reason to expect that our various interpersonal (and intrapersonal) biases will interact. Marie, as we have seen, is prone to mistrust herself, as a result of epistemic injustices accumulated over time, and she thus is at risk of putting too much weight on the views of others, and failing to heed her own best-considered past judgments. Adam, conversely, is overly confident of his own judgments, and so less likely to give sufficient weight to the views of others, and perhaps less likely, too, to attend to the training on implicit bias that he himself once found to be highly credible. In light of such interactions, we should not think of these various cases as merely parallel or analogous. Rather, all these forms of bias feed into a single, overlapping problem: that, from an epistemic point of view, our biases ensure that we do not play well with others.

An adequate understanding of these issues would require coming to terms with the various biases at work, and understanding the ways in which they interact. We would need to be able to assess which of these biases are tolerably accurate, in one or another context, which are likely to distort our judgments, and how these effects accumulate over time. As we come to understand these issues better, the result may be a heightened sense of the injustices that permeate our epistemic lives, but we may also discover that in some ways the situation is less bad than some have supposed. At least for now, to my mind, the appropriate stance is a healthy degree of skepticism regarding our capacities for successful epistemic engagement with others.

## 5 | CONCLUSION

I have argued that our social biases—founded on race, gender, and other categories—should lead us to be skeptical about our ability to assess others accurately. My further argument has been that we should understand such bias skepticism in the context of a broader interpersonal skepticism that arises whenever we make judgments about the worth of others. To be sure, our errors will be a matter of incremental miscalibration rather than radical illusion. Even so, incremental cases accumulate. And inasmuch as knowledge is largely a social phenomenon, and so requires an accurate assessment of others' contributions, doubts about the reliability of our interpersonal epistemic judgments can rightly be said to give rise to a very pervasive form of skepticism.

Some readers will be less pessimistic about whether these biases should drive us all the way to skepticism. They may think our social biases can be kept within known and tolerable limits, or that their impact will be only marginal. In the same way, when it comes to peer disagreement, readers will be variously more or less sympathetic to a skeptical verdict. My ultimate aim, however, is not to make the case for skepticism, but rather to describe the deeper epistemic structure of these questions. Once that is revealed, various hitherto independent parties should be in a position to recognize themselves as working on pieces of the same puzzle.

Moreover, appreciating the shared structure of these various cases allows us to see just how deep the epistemic difficulties run. It is commonly supposed that the way forward out of one or another of these doxastic dilemmas is to rid ourselves of our offending biases. To be sure, we would all like to do that. But when we grasp these issues in their full generality, it becomes clear that there is no straightforward path to epistemic justice through mere moral improvement. In all of these cases, there are distinctively epistemic issues at work. Accordingly, even under the most generously idealizing of assumptions—for instance, that we can recognize our epistemic peers, or the degree of our racial bias—it is surprisingly unclear what the proper path forward is. This is

not a plea for excuses on behalf of the various biases we carry around. The whole point of thinking hard about these issues is to try to do better. Yet, frustratingly, it is not entirely clear, at present, what doing better involves.<sup>23</sup>

## ENDNOTES

- <sup>1</sup> Fricker (2007), p. 4. For the current state of the art regarding epistemic injustice, see Kidd, Medina, and Pohlhaus (2017).
- <sup>2</sup> Other forms of prejudice might of course get added to the list in various cultural contexts, and their enumeration is important to the larger project this paper describes. See, for instance, Peters (2019), which focuses on ideological bias in academia. Here I will focus on race and gender.
- <sup>3</sup> Gournay (1998), p. 35. I have not found other discussions of this passage, but for brief remarks about Gournay's historical context, see Shapiro (2016). For a still earlier example, see Christine de Pizan (1982), which offers an extended defense, circa 1400, of the epistemic and moral credibility of women.
- <sup>4</sup> On this tendency in the literature, see Dotson (2014), who targets the "assumption that epistemic forms of oppression are generally reducible to social and political forms of oppression" (p. 116). Still, as the references throughout this paper should make clear, the epistemic dimension of these issues has by no means been neglected entirely, and extends beyond discussions of epistemic injustice to include, for instance, the epistemologies of ignorance and still broader theoretical work on race and gender (see, e.g., the papers collected in Sullivan and Tuana (2007). Engagement with some of the best work in the field (e.g., Lugones, 2003) requires venturing beyond the narrow boundaries of academic philosophy conservatively defined.
- <sup>5</sup> Subsequent references to Saul (2013) in the main text are by page number alone, within parentheses.
- <sup>6</sup> See Brownstein (2018) ch. 7, Madva (2017), and the papers collected in Sherman and Goguen (2019). For classroom strategies, see Daoust et al. (n. d.).
- <sup>7</sup> For brief recent summaries of the case for implicit bias's significance, see the Introduction to Brownstein and Saul (2016), as well as Payne, Niemi, and Doris (2018). For critical discussions within psychology, see Mitchell and Tetlock (2017) and Singal (2017). For a recent scholarly response, see Jost (2019). Another source of doubt from within psychology over the epistemic significance of bias is the research of Lee Jussim, who concludes that "the power of expectations to distort social beliefs through biases and to create actual social reality through self-fulfilling prophecies is, in general, so small, fragile, and fleeting that it is quite difficult to make a convincing case based on a complete and careful reading of the actual scientific data that such effects likely constitute a major source of inequality" (2012, p. 422). Jussim's work has been largely ignored by philosophers working on epistemic injustice, but poses a direct challenge to many claims in this area.
- <sup>8</sup> For a wide construal of Saul's conclusion, see Director (2018).
- <sup>9</sup> On the usefulness of biases, see the classic discussion in Gigerenzer and Goldstein (1996). A response to bias skepticism along these lines has been made by Antony (2016). Although she accepts the core of Saul's concerns about the seriousness of social bias, she takes a Quinean naturalized attitude to the problem, regarding bias as an omnipresent feature of all cognition rather than a special reason for skeptical despair. Although I will not specifically take up Antony's challenge, the upshot of my discussion will be that special features of interpersonal bias create particularly serious difficulties in these domains. The methods of naturalized epistemology might make a welcomed contribution here, in diagnosing the mechanisms in play, precisely because the skeptical threats are so serious.
- <sup>10</sup> This claim has been defended in detail by Keller (2004), (2018), Paul and Morton (2018a), and Stroud (2006). For a discussion that extends beyond the case of belief, see Wallace (2012). For a duty to believe between persons more generally, see Marušić and White (2018). Whether or not the grounds for these doxastic duties are epistemic is a contentious matter on which here I take no stand. Hence this line of thought is consistent with the arguments of those who insist that the norms of friendship are consistent with evidentialism. For this view see Arpaly and Brinkerhoff (2018), Hawley (2014), and Kwall (2013).
- <sup>11</sup> This is, at any rate, the standard view in the literature on various forms of pragmatic and moral encroachment on belief. For a recent unorthodox argument that such considerations should affect our credences as well as our beliefs, see Gao (2019).
- <sup>12</sup> I owe this line of thought to an anonymous referee.

- <sup>13</sup> See e.g. Bolinger (2020), Enoch, Spectre, and Fisher (2012), and Redmayne (2008). It is contentious whether the problem with statistical evidence is narrowly epistemic or instead violates some extra-evidential moral norm. For present purposes it is enough simply that there is *something* problematic about statistical evidence in cases such as this. Assessing the issue fully would require identifying the problem precisely and then assessing whether recalibration as described is problematic in exactly that way.
- <sup>14</sup> I owe this line of thought to an anonymous referee. Even using statistical information to shape inquiry can be problematic, however, as when it amounts to racial profiling (Lever, 2017). It seems unproblematic, for instance, to give special scrutiny to concert attendees, in light of the information that most of them did not hold tickets. But would members of the above-described fraternity be right to complain if Professor White gave special scrutiny to their papers?
- <sup>15</sup> The impact of “cumulative” and “persistent” epistemic injustices is a reoccurring theme in Fricker (2007), e.g., pp. 21, 46, 54–55. For empirical doubts over the extent to which the impacts of bias are cumulative, see Jussim (2012), ch. 14. For a recent discussion of polarization in interpersonal judgment, see Spaulding (2018), who attempts to account for the phenomenon through the way bias distorts our capacity for social interpretation (or “mindreading”).
- <sup>16</sup> We might say that Adam is, or at any rate wants to be, a feminist “ally.” For doubts over the role of allies, see Brown and Ostrove (2013) and McKinnon (2017). The present remarks perhaps shed some light on why reliable allies can seem so hard to find.
- <sup>17</sup> For the sort of ignorance that can arise in such cases (focusing on racial perspective), see Mills (2007) and, more recently, Medina (2016). On the importance of first-hand experience in cases such as this, see Krishnamurthy (in progress). She argues that, without first-hand experience of racism, it is difficult not just to understand the phenomenon but also, critically, to acquire the motivation to take action against it.
- <sup>18</sup> For the “runaway” collapse of credibility that can occur in cases with this sort of structure, see Jones (2002), pp. 159–160. In this connection, gaslighting can be understood as a pattern of behavior (intended or not) that facilitates this sort of collapse. See Abramson (2014). As with epistemic injustice in general, one might urge that gaslighting involves epistemic complexities that run beyond current discussions of the phenomenon from a purely sociopolitical perspective.
- <sup>19</sup> Closely connected to these risks is the phenomenon of stereotype threat. For a recent philosophical discussion of the cognitive mechanisms at work see Mallon (2016). In such cases it becomes possible that stereotypes become a self-fulfilling prophecy, when people take on the biases (negative and positive) that are endemic to the culture they inhabit. Mallon quotes Appiah (2005), p. 66: “ideas [about social identity] shape the ways people conceive of themselves and their projects.” Here lies a very bleak solution to bias skepticism, where our biases shape how the world actually is.
- <sup>20</sup> Worries over skepticism have motivated the disagreement literature from the start—see, e.g., the final section of Christensen (2007). It is worth noting, in passing, that the comparison to bias skepticism suggests a friendly amendment to the standard skeptical argument from disagreement. That argument runs through an initial *confidence* in the correctness of conciliation. This, as has been widely noted, immediately risks self-refutation, since the conciliatory view is itself subject to considerable disagreement among experts. The argument might be reformulated, however, in line with bias skepticism, so that its starting point is an initial *uncertainty* over what weight to give the judgment of one’s interlocutor. As with bias skepticism, one would accordingly have no good basis for any sort of recalibration. Such uncertainty yields an even more thoroughgoing skepticism, but without the risk of self-refutation, because it captures the meta-level uncertainty over whether one ought to be conciliatory.
- <sup>21</sup> For the view that conciliatory views neglect the evidence, see Kelly (2010) and Enoch (2010). On various aspects of self-trust, see Foley (2001), Govier (1998), Paul and Morton (2018b), and Preston-Roedder (2018). On self-trust in the context of peer disagreement see Pasnau (2015).
- <sup>22</sup> For an extended argument in favor of privileging the present self, see Foley (2001), p. 149: “it cannot be a demand of rationality that I shed my current perspective and adopt a vantage point from which I treat all of my temporal selves and their opinions identically.” For the question of whether there are special intrapersonal demands of rationality that hold across time, see Heddon (2015). For further discussion of our bias toward the present self, see Pasnau (2017), lecture 5.
- <sup>23</sup> I could not have written this paper without the help I have received from many interlocutors, over many iterations of these ideas, including David Boonin, Heather Demarest, Chris Heathwood, Mitzi Lee, Colin Macleod, Andrei Marasoiu, Elliot Paul, Robert Rupert, Raul Saucedo, Julia Staffel, Brian Talbot, Jennifer Wang, and

audiences at the University of Victoria (BC), the University of Virginia, the University of St. Andrews, and the University of Colorado. Extremely useful reports from a series of anonymous readers have made the final paper much better.

## REFERENCES

- Abramson, K. (2014). Turning up the lights on gaslighting. *Philosophical Perspectives*, 28, 1–30.
- Anderson, E. (2012). Epistemic justice as a virtue of social institutions. *Social Epistemology*, 26, 163–173.
- Antony, L. (2016). Bias: Friend or foe? Reflections on Saulish skepticism. In M. Brownstein & J. Saul (Eds.), *Implicit Bias and Philosophy, Volume 1: Metaphysics and Epistemology* (pp. 157–190). Oxford: Oxford University Press.
- Appiah, K. A. (2005). *The Ethics of Identity*. Princeton, NJ: Princeton University Press.
- Arpaly, N., & Brinkerhoff, A. (2018). Why epistemic partiality is overrated. *Philosophical Topics*, 46, 37–51.
- Basu, R. (2019). The wrongs of racist beliefs. *Philosophical Studies*, 176, 2497–2515.
- Basu, R., & Schroeder, M. (2019). Doxastic wronging. In B. Kim & M. McGrath (Eds.), *Pragmatic Encroachment in Epistemology* (pp. 158–178). London: Routledge.
- Bolinger, R. J. (2020). The rational impermissibility of accepting (some) racial generalizations. *Synthese*, 197, 2415–2431.
- Bramson, A., Grim, P., Singer, D. J., Berger, W. J., Sack, G., Fisher, S., ... Holman, B. (2017). Understanding polarization: Meanings, measures, and model evaluation. *Philosophy of Science*, 84, 115–159.
- Brown, K., & Ostrove, J. (2013). What does it mean to be an ally? The perception of allies from the perspective of people of color. *Journal of Applied Social Psychology*, 43, 2211–2222.
- Brownstein, M. (2018). *The Implicit Mind: Cognitive Architecture, the Self, and Ethics*. Oxford: Oxford University Press.
- Brownstein, M., & Saul, J. (2016). *Implicit Bias and Philosophy* (Vol. 2). Oxford: Oxford University Press.
- Christensen, D. (2007). Epistemology of disagreement: The good news. *The Philosophical Review*, 116, 187–217.
- Christine de Pizan (1982). *The Book of the City of Ladies*, tr. E. J. Richards, (Ed.). New York: Persea Books.
- Daoust, L. et al. (n. d.). Best practices for the inclusive philosophy classroom. Retrieved from <https://phildiversity.weebly.com>
- Director, S. (2018). A dilemma for Saulish skepticism: Either self-defeating or not even skepticism. *Disputatio*, 10, 43–55.
- Dotson, K. (2014). Conceptualizing epistemic oppression. *Social Epistemology*, 28, 115–138.
- Enoch, D. (2010). Not just a truthometer: Taking oneself seriously (but not too seriously) in cases of peer disagreement. *Mind*, 119, 953–997.
- Enoch, D., Spectre, L., & Fisher, T. (2012). Statistical evidence, sensitivity, and the legal value of knowledge. *Philosophy and Public Affairs*, 40, 197–224.
- Feldman, R., & Warfield T. A. (Eds.). (2010). *Disagreement*. Oxford: Oxford University Press.
- Foley, R. (2001). *Intellectual Trust in Oneself and Others*. Cambridge: Cambridge University Press.
- Fricke, M. (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press.
- Fricke, M. (2013). Epistemic justice as a condition of political freedom? *Synthese*, 190, 1317–1332.
- Gao, J. (2019). Credal pragmatism. *Philosophical Studies*, 176, 1595–1617.
- Gendler, T. S. (2011). On the epistemic costs of implicit bias. *Philosophical Studies*, 156, 33–63.
- Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, 103, 650–669.
- Gournay, M. le Jars de (1998). *Preface to the Essays of Michel de Montaigne*, tr. In R. Hillman & C. Quesnel (Eds.). Tempe, AZ: Medieval and Renaissance Texts and Studies.
- Govier, T. (1998). *Dilemmas of Trust*. Montreal: McGill-Queen's University Press.
- Haslanger, S. (2015). Social structure, narrative and explanation. *Canadian Journal of Philosophy*, 45, 1–15.
- Hawley, K. (2014). Partiality and prejudice in trusting. *Synthese*, 191, 2029–2045.
- Heddon, B. (2015). *Reasons without Persons: Rationality, Identity, and Time*. Oxford: Oxford University Press.
- Jones, K. (2002). The politics of credibility. In L. Antony & C. Witt (Eds.), *A Mind of One's Own: Feminist Essays on Reason and Objectivity* (pp. 154–176). Boulder: Westview Press.
- Jones, K. (2012). The politics of intellectual self-trust. *Social Epistemology*, 26, 237–251.

- Jost, J. T. (2019). The IAT Is dead, long live the IAT: Context-sensitive measures of implicit attitudes are indispensable to social and political psychology. *Current Directions in Psychological Science*, 28, 10–19.
- Jussim, L. J. (2012). *Social Perception and Social Reality: Why Accuracy Dominates Bias and Self-Fulfilling Prophecy*. New York: Oxford University Press.
- Kawall, J. (2013). Friendship and epistemic norms. *Philosophical Studies*, 165, 349–370.
- Keller, S. (2004). Friendship and belief. *Philosophical Papers*, 33, 329–351.
- Keller, S. (2018). Belief for someone else's sake. *Philosophical Topics*, 46, 19–35.
- Kelly, D., & Roedder, E. (2008). Racial cognition and the ethics of implicit bias. *Philosophy Compass*, 3, 522–540.
- Kelly, T. (2010). Peer disagreement and higher-order evidence. In R. Feldman & T. A. Warfield (Eds.), *Disagreement* (pp. 111–174). Oxford: Oxford University Press.
- Kidd, I. J., Medina, J., & Pohlhaus Jr., G. (Eds.) (2017). *The Routledge Handbook to Epistemic Injustice*. London: Routledge.
- Krishnamurthy, M. (in progress). White moral blindness.
- Lever, A. (2017). Race and racial profiling. In N. Zack (Ed.), *The Oxford Handbook of Philosophy and Race* (pp. 425–435). New York: Oxford University Press.
- Lugones, M. (2003). *Pilgrimages = Peregrinajes: Theorizing Coalition against Multiple Oppressions*. Lanham, MD: Rowman & Littlefield.
- McKinnon, R. (2017). Allies behaving badly: Gaslighting as epistemic injustice. In I. J. Kidd, J. Medina, & G. Pohlhaus (Eds.), *The Routledge Handbook to Epistemic Injustice* (pp. 167–174). London: Routledge.
- Madva, A. (2017). Biased against debiasing: On the role of self-transformation in the struggle against prejudice. *Ergo*, 4, 145–179.
- Mallon, R. (2016). Stereotype threat and persons. In M. Brownstein & J. Saul (Eds.), *Implicit Bias and Philosophy* (Vol. 1, pp. 130–154). Oxford: Oxford University Press.
- Marušić, B., & White, S. (2018). How can beliefs wrong? A Strawsonian epistemology. *Philosophical Topics*, 46, 97–114.
- Medina, J. (2016). Ignorance and racial insensitivity. In R. Peels & M. Blaauw (Eds.), *The Epistemic Dimensions of Ignorance* (pp. 178–201). Cambridge: Cambridge University Press.
- Mills, C. (2007). White ignorance. In S. Sullivan & N. Tuana (Eds.), *Race and Epistemologies of Ignorance* (pp. 11–38). Albany, NY: State University of New York Press.
- Mitchell, G., & Tetlock, P. E. (2017). Popularity as a poor proxy for utility: The case of implicit prejudice. In S. O. Lilienfeld & I. D. Waldman (Eds.), *Psychological Science under Scrutiny: Recent Challenges and Proposed Solutions* (pp. 164–195). Hoboken, NJ: Wiley Blackwell.
- Pace, M. (2011). The epistemic value of moral considerations: Justification, moral encroachment, and James' "Will to believe". *Nous*, 45, 239–268.
- Pasnau, R. (2015). Disagreement and the value of self-trust. *Philosophical Studies*, 172, 2315–2339.
- Pasnau, R. (2017). *After Certainty: Historical Reflections on Our Epistemic Ideals and Illusions*. Oxford: Oxford University Press.
- Pasnau, R. (2018). Belief in a fallen world. *Res Philosophica*, 95, 531–559.
- Paul, S. K., & Morton, J. M. (2018a). Believing in others. *Philosophical Topics*, 46, 75–95.
- Paul, S. K., & Morton, J. M. (2018b). Grit. *Ethics*, 129, 175–203.
- Payne, K., Niemi, L., & Doris, J. M. (2018). How to think about 'implicit bias'. *Scientific American*, March 27, 2018. Retrieved from [www.scientificamerican.com/article/how-to-think-about-implicit-bias](http://www.scientificamerican.com/article/how-to-think-about-implicit-bias)
- Peters, U. (2019). Implicit bias, ideological bias, and epistemic risks in philosophy. *Mind and Language*, 34, 393–419.
- Preston-Roedder, R. (2018). Three varieties of faith. *Philosophical Topics*, 46, 173–199.
- Redmayne, M. (2008). Exploring the proof paradoxes. *Legal Theory*, 14, 281–309.
- Saul, J. (2013). Scepticism and implicit bias. *Disputatio*, 5, 243–263.
- Shapiro, L. (2016). Revisiting the early modern philosophical canon. *Journal of the American Philosophical Association*, 2, 365–383.
- Sherman, B. R., & Goguen, S. (Eds.) (2019). *Overcoming Epistemic Injustice: Social and Psychological Perspectives*. Lanham, MD: Rowman & Littlefield.
- Singal, J. (2017). Psychology's favorite tool for measuring racism isn't up to the job. *New York Magazine*, Retrieved from <https://www.thecut.com/2017/01/psychologys-racism-measuring-tool-isnt-up-to-the-job.html>

- Singer, D. J., Bramson, A., Grim, P., Holman, B., Jung, J., Kovaka, K., ... Berger, W. J. (2019). Rational social and political polarization. *Philosophical Studies*, 176, 2243–2267.
- Spaulding, S. (2018). Do you see what I see? How social differences influence mindreading. *Synthese*, 195, 4009–4030.
- Stroud, S. (2006). Epistemic partiality in friendship. *Ethics*, 116, 498–524.
- Sullivan, S., & Tuana, N. (Eds.) (2007). *Race and Epistemologies of Ignorance* (Vol. 2007). Albany: SUNY Press.
- Wallace, R. J. (2012). Duties of love. *Proceedings of the Aristotelian Society*, suppl. vol. 86, 175–198.
- Zheng, R. (2016). Attributability, accountability, and implicit bias. In M. Brownstein & J. Saul (Eds.), *Implicit Bias and Philosophy, Volume 2: Moral Responsibility, Structural Injustice, and Ethics* (pp. 62–89). Oxford: Oxford University Press.

**How to cite this article:** Pasnau R. Bias and interpersonal skepticism. *Noûs*. 2020;1–22.  
<https://doi.org/10.1111/nous.12352>