

Profiling, Neutrality, and Social Equality

Lewis Ross

To cite this article: Lewis Ross (2022) Profiling, Neutrality, and Social Equality, Australasian Journal of Philosophy, 100:4, 808-824, DOI: [10.1080/00048402.2021.1926522](https://doi.org/10.1080/00048402.2021.1926522)

To link to this article: <https://doi.org/10.1080/00048402.2021.1926522>



© 2021 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group



Published online: 08 Jun 2021.



Submit your article to this journal [↗](#)



Article views: 1915



View related articles [↗](#)



View Crossmark data [↗](#)



Profiling, Neutrality, and Social Equality

Lewis Ross

London School of Economics and Political Science

ABSTRACT

I argue that traditional views on which beliefs are subject only to purely epistemic assessment can reject demographic profiling, even when based on seemingly robust evidence. This is because the moral failures involved in demographic profiling can be located in the decision not to suspend judgment, rather than supposing that beliefs themselves are a *locus* of moral evaluation. A key moral reason to suspend judgment when faced with adverse demographic evidence is to promote social equality—this explains why positive profiling is dubious, along with more familiar cases of negative profiling, and why profiling is suspect even when no particular action is at stake. My suspension-based view, while compatible with revisionary normative positions, does not presuppose them. Philosophers of all stripes can reject demographic profiling both in thought and deed.

ARTICLE HISTORY Received 12 August 2020; Revised 13 April 2021

KEYWORDS profiling; moral encroachment; social egalitarianism; statistical evidence; suspended judgment; neutrality; ethics of belief

1. Introduction

Demographic profiling is philosophically puzzling. Those who engage in profiling—drawing inferences about individuals by relying on statistical generalisations about their sex, race, or socioeconomic class—are often open to criticism even when relying upon robust evidence. To explain why, some suggest that we must accept revisionary theories on which beliefs are susceptible to moral evaluation in hitherto-unexpected ways. This paper argues that we can reject profiling, both in thought and deed, without endorsing revisionary commitments that blur epistemic and moral evaluation of belief.

The core argument is that, by closely attending to the nature of the neutral attitudes, particularly of suspended judgment, we can maintain that beliefs should only be assessed against purely epistemic standards, even if there can be moral reasons to suspend judgment prior to forming a settled view. Drawing on social egalitarian theory, I aim to show why beliefs that result from profiling are intrinsically undesirable—regardless of whether they lead to any particular action. Appreciating the intrinsic value of social equality allows us to explain why *positive* profiling is often deleterious, along with more familiar cases of negative profiling. Recognising the importance of social equality and the risks involved in profiling, so I claim, provides

reason to adopt a default policy of suspending judgment when all that we have available is statistical evidence pertaining to the prevalence of certain virtues, vices, or intelligence-level within some demographic group.

2. Relying on Statistical Evidence

Statistical evidence is criticised from various quarters. Epistemologists have long been preoccupied (chiefly through ‘lottery cases’) with epistemic deficiencies in statistical evidence,¹ philosophers of science have worried about the reference-class problem for probabilistic inferences,² and scholars of jurisprudence have questioned the legal value of naked statistics.³ Despite such concerns, there are clearly cases where bare statistical evidence rationally requires (or at least licenses) belief and action. Take the following case:

Blow a Gasket. You read a report stating that model *M* cars have a marked propensity to break down. Specifically, 20% of them have a fault causing them to blow a gasket within five years. This, so the report says, makes them the least reliable car released in the past decade. At the showroom, you see a selection of cars. You think to yourself that the model *M* car is comparatively less reliable than the other cars. On that basis, you choose to buy another type of car.

While the evidence does not support the probable unreliability of any particular car, it licenses the thought that a given model *M* car is *probably less reliable* than those belonging to other models. It seems rational to act on the basis of that comparative belief by choosing not to purchase a model *M* car.

Perhaps there are special prohibitions about relying on unfavourable statistical evidence about human beings? Cases inspired by debates about the proof paradox suggest not. Adverse statistical evidence does seem to license various cognitive and practical responses in the interpersonal domain, even if not in the courtroom. Take a case (lightly adapted) from the legal literature:

Attackers. There are 20 employees on a cruise ship working on deck. 19 of them attack, rob, and sadistically torture a wealthy passenger, putting into action a plan about which the 20th employee knew nothing. The 20th employee played no role in the assault and could have done nothing to stop it. There is no further evidence available to settle the question of any particular employee’s involvement.⁴

Suppose that, as should happen according to many legal philosophers, all employees are released without criminal charge because the evidence is insufficient for conviction.⁵ Now suppose that one of these released employees moves in next-door to you. You have read about the case and know all about the incriminating statistical evidence. It is arguably not permissible to act on the bald proposition *my neighbour is dangerous*—for instance, by publicly denouncing them.⁶ However, it nonetheless seems epistemically rational to form the belief that, given your evidence, there is a

¹ Hawthorne [2004] provides wide-ranging discussion. For overview and empirical study, see Ebert et. al [2018].

² Reichenbach [1949] provides a canonical statement. Hajek [2007] provides further discussion.

³ For overview, see Redmayne [2008] or Ross [2020].

⁴ This scenario has been adapted from the famous Prisoners case in Redmayne [2008], where the potential attackers are prisoners in a yard.

⁵ See Enoch et. al [2012], Smith [2018], Pritchard [2018], Gardiner [2019], and Littlejohn [2020] for a selection of arguments for the conclusion that legal verdicts should not be based on naked statistics.

⁶ Buchak [2014] discusses the role of outright belief in legitimising blame.

heightened chance of your neighbour being dangerous. This is supported by the fact that it is prudentially rational to modify your behaviour.⁷ For instance, it would be reasonable to initially adopt a heightened degree of caution in your dealings with them—for example, by refraining from giving them your spare-key for safe keeping, as you would usually do with a new arrival to your sleepy town.⁸

It appears that we can rationally form comparative beliefs about humans by using adverse statistical evidence, and we can moderate our behaviour accordingly. What, then, is special about cases of demographic profiling? It will help to focus on an example.⁹

Diner. You work at a diner. From time to time, teaspoons are stolen. You know, based on a recent extensive analysis of CCTV footage, that members of demographic Y are responsible for 90% of all teaspoon thefts. Suppose that demographic Y is distinctively marked out by a certain appearance. One morning, you wait on only two tables: one with members of demographic Y and another with members of demographic Z. Taking their collective dishes to be washed, you notice that a number of teaspoons have disappeared.

Two key questions arise:

The practical question: how should one act, in light of adverse demographic evidence?

The cognitive question: what is the best cognitive response to adverse demographic evidence?

Let's start with the practical question. Ought you to act on the basis of the adverse demographic evidence—say, by pursuing the members of Y and accusing them of theft, or by treating them with heightened suspicion when serving them? It seems compelling to suppose not. But why? What is special about demographic statistics that distinguishes them from the aforementioned cases?

We already possess the outline of an answer to the practical question. Legal scholars have noted that relying upon adverse demographic evidence exposes members of certain groups to a higher *ex ante* risk of mistaken conviction, compared to other defendants facing similar charges [Di Bello and O'Neill 2020]. For, by making certain demographics susceptible to inculpatory evidence not used against other groups, they would find themselves convicted of crimes more often. Even if we grant that profile evidence is fairly reliable, it is not infallible. Given that mistakes invariably occur when using any type of fallible evidence, the result would be that members of the profiled group would suffer from more mistaken convictions than others would. The worry is that imposing an increased *ex ante* risk of mistaken conviction on members of certain groups fails to treat them as equals before the law.¹⁰

The broad concern generalises to the interpersonal domain. Since actions performed on the basis of profiling usually involve unwanted consequences for those being profiled, the same thought about raising the risk of erroneously penalising certain people applies outside, as well as inside, the courtroom. Especially when

⁷ I consider the relevance of this fact in relation to legal verdicts elsewhere (see Ross [2021]).

⁸ See Munton [2019] for a useful attempt to identify a purely epistemic flaw in such evidence.

⁹ I'm interested in drawing inferences about individuals from a body of statistical evidence. There are separate issues concerning *generic* statements (e.g. 'sharks attack bathers'). See Leslie [2017].

¹⁰ There are various relevant proposals in the literature on the proof paradox discussing why courts should not base verdicts on mere statistics: for instance, see Wasserman's [1992] focus on autonomy, or Nunn [2015] on due process. See Ross [2020] or Backes [2020: sec. 2.1] for general discussion.

demographic membership is marked by a visually salient sortal characteristic like dress or skin-colour, profiling will lead to members of that demographic being repeatedly exposed to accusations, defensive behaviour, and suspicion. Again, given the fallibility of profile evidence, this unwanted treatment will not always be warranted by the actual behaviour of persons being profiled. While these mistakes are individually unfortunate, their repetition risks instantiating an *oppressive* pattern against members of profiled groups. Andreas Mogensen [2019: 17] draws on Marilyn Frye's evocative 'birdcage analogy' in precisely this context:

just as a single wire is no significant impediment to freedom apart from its place within a network of others, so individual elements of an oppressive society may not appear as morally troubling as they are unless viewed as part of the larger pattern of similar injustices of which they form part.

The focus on oppression also distinguishes demographic profiling from the *Attackers* case, in which behaviour modification in response to statistics was reasonable: in that case, we were relying upon identity-contingent evidence (presence on a cruise-ship at time *t*) that does not risk oppressing any group marked by a persistent salient sortal characteristic. These considerations undergird a powerful reason not to *act* on the basis of adverse demographic evidence.

This leaves the *cognitive question*. Is it permissible to *believe* that an individual member of demographic Y is comparatively more likely than other customers to be a thief, just so long as these beliefs are not made overt?

An affirmative answer (call this the *split response*) is, on its face, compatible with acknowledging downstream harms of profiling. Moreover, the split response comports with some aspects of ordinary practice: there are many beliefs permissible to hold, but impermissible to make manifest that you hold. For instance, one paradigmatic way that we make our beliefs manifest to others is by asserting their contents. There are many examples of rational beliefs that we ought not to assert. Suppose that you know, on impeccable evidence, that a colleague is being passed over for promotion. There might be factors that make asserting this belief impermissible—for example, prudentially impermissible (if you tell them, they will shoot the messenger), or morally impermissible (if you tell them, it will precipitate a breakdown, minutes before their keynote address). The fact that we should keep a belief covert is not necessarily a ground for supposing that we should not hold the belief in the first place.

I will reject the split response, by providing a new rationale for the idea that profiling beliefs, even when covertly held, are intrinsically undesirable. However, before outlining my own positive view, I want to discuss and raise challenges for two prominent alternative attempts to reject the legitimacy of profiling beliefs—the *moral wronging* view and the *moral encroachment* view.

3. Immoral Beliefs and Moral Encroachment

This section discusses two influential theories that reject many profiling beliefs. Both approaches contain insights but have distinctive theoretical costs. The first is that profiling beliefs *morally wrong* their subjects, and the second is that the associated moral risks render profiling beliefs *epistemically irrational*.

3.1 Are Profiling Beliefs a Locus of Moral Evaluation?

There is a natural view about normativity on which the epistemic and the moral are evaluative domains with different objects. According to this conception, our *beliefs* should be assessed against epistemic norms, while our *actions* (plus omissions) and *character* fall under the governance of moral norms. In an influential series of papers, Rima Basu has argued that profiling beliefs provide decisive reason to reject this traditional normative account.¹¹

According to Basu's theory, beliefs themselves are a proper *locus* of moral evaluation. Granting this allows us to say that profiling beliefs *morally wrong* their subjects. (Henceforth, I'll call this *the moral wronging view*). Under the moral wronging view—and, for Basu, this is a key virtue—the moral wrong of profiling is compatible with there being robust evidence rationalising the profiling belief. The central example that she uses to motivate the view is that of the 'rational racist' being criticisable for believing that a black diner will tip less, even if they are aware of empirical studies supporting just this hypothesis [Basu 2019a: 2498]. The clear advantage of the moral wronging view is its providing a straightforward vindication of the idea that the cognitive practices of the 'rational racist' are morally unsavoury.

Issues surrounding doxastic voluntarism are especially pertinent for the moral wronging view, given that it is *prima facie* attractive only to criticise agents for matters within their control. The familiar worry is—in contrast with action—that we cannot believe or disbelieve at will. In one paper, Basu [ibid.: 2514] says that we must choose between either 'giving the rational racist a pass' in virtue of holding that we can only be morally evaluated for that which is under our control (a 'control condition'), or we must deem concerns about control a 'theoretical technicality' that must be abandoned to vindicate the conclusion that racist beliefs wrong their subjects. In later work, Basu and Schroeder go further by mounting a two-pronged attack against control-based objections: first, they advert to various theoretical arguments against control conditions [2019: 188]; second, they suggest that control conditions fail to account for the fact that how we settle a line of inquiry is often 'a matter of our own psychology, and thus up to us' [ibid.: 194].

Although Basu and Schroeder perhaps overstate the extent to which our psychology is in fact up to us, I agree that there is much to recommend rejecting excessively strict control conditions on moral responsibility. However, what we lack is an account of what *positive* response is both possible—and required—when facing a question about which we have only adverse demographic evidence. I take it that it is a natural *desideratum*, one that I further defend in the next paragraph, that we should prefer a solution to the problem of profiling that does not require that we actively believe *against* the evidence. But nor can we simply become like an inanimate object and entertain no cognitive response whatsoever. In the final section, I will aim to outline just such a solution to this issue, drawing on the positive epistemic state of suspended judgment.

Before moving on, it will be dialectically helpful to emphasise two further points regarding the moral wronging view. First, it bears emphasis that treating beliefs as a *locus* of moral wrong is theoretically heterodox. A common view holds that truth (or knowledge) is the 'aim of belief' and that norms concerning belief's evaluation

¹¹ See Basu [2019a, 2019b] and Basu and Schroeder [2019].

concern fundamentally veritistic notions such as strength of evidence or reliability.¹² Familiar extant views that *do* talk about moral duties to have certain beliefs—such as the Cliffordian view that we have a moral duty to believe the truth¹³—are fundamentally veritistic in nature. Of course, a view's being revisionary is no reason to dismiss it, especially not if it manages to capture important data-points surrounding the intuitive wrong of profiling. Nonetheless, preferring the traditional normative dichotomy between the epistemic evaluation of belief and the moral evaluation of action, omission, and character is not merely a conservative preference. A second reason to prefer that we restrict beliefs to epistemic evaluation—and reject views that require, morally speaking, that we believe against the evidence—is to preserve *coordination* between normative domains. What I mean by 'coordination' is the preclusion of certain types of normative conflict. *Prima facie*, it would be preferable to avoid situations in which our best epistemic theory licenses believing *p* but our best moral theory states that believing *p* is morally impermissible.¹⁴ Basu and Schroeder [2019] accept the force of the coordination problem. Their solution is to suggest that this issue might be dealt with by accepting a further view—namely, that epistemic rationality is itself affected by moral considerations. It is to that view that I turn next.

3.2 Do Moral Factors Make Profiling Beliefs Epistemically Irrational?

Although treating belief as a *locus* of moral evaluation is fairly radical, there is another way to incorporate moral considerations into doxastic assessment. This view draws upon an approach in epistemology that argues that pragmatic factors, in addition to truth-related factors, are relevant for epistemic rationality. As the epistemological terminology puts it, the suggestion is that pragmatic factors *encroach* upon epistemic assessment of belief. So, for instance, a now-venerable argument contends that the importance of a belief's truth for you can determine how much evidence is needed for it to count as justified for, or known by, you.¹⁵ Some have sought to extend this idea so as to argue that belief is subject to *moral encroachment*. Encroached views of epistemic rationality remain controversial: many are the arguments both for, and against, pragmatic encroachment. Equally, it is an open question whether pragmatic and moral encroachment stand or fall together (see Fritz [2017] and Bolinger [2020b] for opposing perspectives).

Bolinger [2020a] has a particularly helpful way of framing the moral encroachment response to the profiling problem: she claims that epistemic justification is impinged upon by the *moral risks* that we run by accepting some proposition *p*.¹⁶ On this account, the riskier a proposition, the higher the epistemic standards that we must meet in order to justifiably accept it. Take the *Diner* case: by accepting that one particular party is more likely than the other to be responsible for the theft, we run a

¹² Chan [2013] is a classic collection on the aim of belief.

¹³ See Clifford [1877]. A famous exception comes from William James (see Pace [2011] for contemporary discussion).

¹⁴ Whether this would generate *dilemmas* depends on whether we can speak sensibly about the correct response 'all things considered'.

¹⁵ See DeRose [1992] for the best-known example.

¹⁶ Bolinger remains neutral on the relationship between belief and acceptance, noting that they are identical on some views but not others. Alternative versions of moral encroachment focus on knowledge: see, for example, Moss's [2018] treatment. An important predecessor is Pace [2011]; Gardiner [2018] provides helpful discussion. Jackson and Fritz [forthcoming] discuss moral encroachment on credences.

heightened risk of violating their interests by exposing them to wrongful accusation, a risk underscored by the fact that they have a heightened susceptibility to profiling because of a visible sortal feature that they possess. Given these risks, so the thought goes, the epistemic standards required for justified acceptance are raised beyond the point at which they are met by the merely statistical evidence that we have available.

Notably, the moral encroachment view avoids the normative coordination problem because it claims that moral factors only play a role in *determining whether a belief is epistemically rational*. Beliefs are not treated as a *locus* of moral evaluation in their own right. Another attractive feature of this view is its explaining why profiling is problematic even if we happen to believe *truly* that the subject of the profile possesses some negative characteristic. Bolinger writes thus [ibid: 2423, emphasis added]:

Having accepted that *p*, *S* might directly harm *A* by acting as if *p*. But even if she luckily avoids imposing these harms—either because *p* is in fact true, or because *S* never has the opportunity to act on her assumption that *p*—she wrongs *A* if she accepts *p* without adequate justification. She’s closed inquiry too early, *recklessly exposing A to unjustified risk of suffering the harms involved in p-based mistakes*.

However, like the moral wronging view, simply arguing for the morally encroached nature of epistemic rationality leaves us without a *positive* characterisation—that doesn’t involve believing against the evidence—of how we should respond to negative profiles. Moreover, we might have some doubts about the scope and nature of the prohibitions generated by the moral encroachment view. A first is that the moral encroachment view diverges from the moral wronging view in how it diagnoses what I’ll call *inert profiling*—cases in which acting on the basis of profiling beliefs is precluded. These cases come in two different varieties. One type of inert profiling is when no action is at stake. For example, suppose that someone is sitting at home, idly flicking through a magazine, forming profiling beliefs about those photographed therein. A second type of inert profiling is when there *is* a profiling-relevant action at issue, but the agent has, for some independent reason, committed not to act. Basu provides an instructive example: suppose that a security guard has the job of preventing shoplifting, they profile someone as a shoplifter, but they are committed to doing nothing because they hate the company for which they work. Intuitions may vary, but I agree with Basu that inert profiling is *prima facie* objectionable. The moral wronging view says that the beliefs involved in inert profiling are wrong. It is not obvious how the moral encroachment view—which ties epistemic rationality to the moral risks of acting on beliefs—can explain discomfort with profiling when acting impermissibly is ruled out by circumstance or prior commitment. A second related issue is that any prohibition on profiling under a view that focuses on expected moral costs must be *contingent*. This is because the agent is enjoined to weigh the strength of the evidence against the expected costs of having the belief. There are reasons to worry about this sort of approach. One reason is that we can arbitrarily increase the strength of the evidence supporting the negative profile or minimise the expected moral costs yet plausibly still have intuitive resistance to such profiling, while another, perhaps deeper, worry is that we might expect the morally responsible agent to have a more fundamental aversion to profiling, rather than its being contingent upon weighing and balancing the costs in each particular case. A third issue—relevant to both the moral wronging and moral encroachment view—is that these approaches don’t straightforwardly explain what is wrong with *positive* profiling (for

example, automatically supposing that an Asian person is good at mathematics). It is, I think, an open question whether we always *wrong* the subject of a positive profile. Suppose that someone receives some enviable position due to positive profiling. While they have not been treated as an equal, it does not strike me as obvious that this *necessarily* harms the subject of the profile. Similarly, for moral encroachment views, it is not clear that there will be strong moral risks associated with every case where we might positively profile somebody.

I don't claim that these are insurmountable problems, although they are important data-points that I think we ought to capture. I begin the defence of my own view in the next section.

4. Profiling and Social Equality

I'll begin by highlighting a purely general problem caused by profiling beliefs, not contingent on some identifiable wrong or harm being done to any individual. In doing so, I will draw upon a particular conception of justice—namely, social egalitarianism.

Departing from G.A. Cohen's [1989: 906] adage that justice has some 'currency' that ought to be distributed fairly, social (or relational) egalitarians stress the non-material components of equality.¹⁷ Take class differences, such as aristocracies. While class differences are often accompanied by distributive inequalities—money, manor houses, etc.—what is objectionable about such hierarchies is not reducible to distributive inequality. Some of these problems might be explained with reference to how they inhibit other valuable ideals, such as the ideals of democratic governance. However, some of what is problematic about the existence of such hierarchies is simply *attitudinal*. It is objectionable for a society to believe that there are people of inherently superior stock, that there is a class of betters or elites, even if these statuses did not entail an unequal distribution of material goods or opportunities. Social equality does not, however, require that we form no opinions about the strengths and weaknesses of our fellow citizens. A standard distinction is often drawn between hierarchies of *status*, such as those holistic evaluative distinctions found in aristocratic or caste-based societies, and hierarchies of *esteem*, which pertain to specific characteristics.¹⁸ Although hierarchies of status are generally objectionable, social equality is not threatened by *some* differences in esteem in so far as these can be legitimately earned. For example, experts who earn intellectual deference because of their good judgment are compatible with social egalitarian ideals. What relational equality *does* require is a level playing field with respect to earning the esteem of your fellow citizens.

This level playing field obviously requires material goods. One is hard-pressed to earn esteem as a scientist if denied a rudimentary education. However, bearing in mind my suggestion that relational (in)equality has an attitudinal dimension, there are also constraints placed upon the cognitive attitudes that it is desirable to possess. Specifically, I think that relational equality presupposes a *default attitude of indifference* about whether any particular person possesses certain characteristics deserving of heightened or lowered esteem. The esteem-worthy characteristics about which we

¹⁷ See Fourie et al. [2015] for a collection of such writings. The roots of analytic social egalitarianism were developed in a body of work by, among others, Iris Marion Young [1990], David Miller [1997], Jonathan Wolff [1998], and Elizabeth Anderson [1999]. Wolff [2019] provides historical discussion. The precise relationship between distributive and social equality is contested: see Scheffler [2015] and Lippert-Rasmussen [2018].

¹⁸ For example, see Fourie et al. [2015] or Wolff [2019] for discussion.

should have this default indifference includes, at a minimum, intelligence, virtue, and vice, but there is also an argument—which I won't pursue here—that it might also include various domain-specific skills and competencies (for example, technical or sporting abilities). When I refer to an attitude of indifference, I do not mean indifference in the sense of not caring, but rather in the stricter etymological sense of not making a distinction. This indifference is a *default* because it is what we should assume until we have personalised evidence that disrupts our indifference. The idea is simple: prior to someone distinguishing themselves in one way or another, we should not believe them to be more or less deserving of esteem than anybody else is.¹⁹

The beliefs that result from demographic profiling undermine the attitudinal requirements of social equality. By harbouring antecedent beliefs about the esteem-relevant characteristics of our fellows before an individual has had the opportunity to personally distinguish themselves in one way or the other, we are not providing those whom we encounter with a level playing field. Not only are such attitudes at odds with social equality in themselves, but they also have deleterious downstream cognitive effects. Attitudes of social equality are important because they facilitate what we can term *evidential parity*. Evidential parity has two parts—one concerning positive assessment, and the other concerning negative assessment. Social equality demands that we ought not (i) to give members of certain groups a heightened benefit of the doubt, or (ii) to tend to interpret evidence as confirming the worst about certain groups over others. Having beliefs characteristic of profiling is in tension with these dispositions. Take a straightforward example. Suppose that you believe that a member of some demographic X is, in virtue of their demographic, more likely to be a bad driver. Suppose that you see them perform a careless manoeuvre. You will be primed to see this evidence as confirming your antecedent belief about their poor driving, rather than as an exception, or as an isolated event with scant evidential value. You would have the opposite reaction to the same evidence about someone whom you profile as a cautious driver. In this sense, your prior cognitive attitudes will lead you to interpret new evidence in fundamentally different ways. This draws out the connection between social equality and another valuable good—namely, autonomy. The cognitive components of social equality are important because they facilitate the *autonomy* of those whom we encounter: it enables them to self-author how they are received by their fellow citizens, rather than to have the reception of their behaviour coloured by prior assumptions.²⁰ The foregoing considerations, I suggest, also provide an instructive diagnosis of why positive profiling is attitudinally undesirable even if it does not involve any harm to the person(s) subject to the positive profile. Just as according someone a lofty status—such as membership of a privileged caste—offends social equality, so does according someone a default assumption of heightened esteem.

Of course, agreeing that there are downstream cognitive and practical harms that typically stem from profiling beliefs is separable from being persuaded that the beliefs themselves are *intrinsically* destructive of social equality. Although making this point is perhaps not strictly required for my overall argument, I am in fact

¹⁹ See Liao and Ye [forthcoming] for a recent attempt to delineate the boundary between individualised and statistical evidence.

²⁰ My argument here could be viewed a sympathetic development of Wasserman's [1992] work on autonomy and the proof paradox: see Pundik [2008] for discussion. See Garrau and Laborde [2015] for further discussion of social equality and autonomy.

persuaded by the stronger view: profiling attitudes are intrinsically inconsistent with the ideals of social equality. We have already discussed some suggestive cases in favour of this conclusion—namely, *inert profiling* where there is no particular action under consideration, and *positive profiling* where risks to the subject of the profile might not be apparent. These, I suggest, are undesirable in their own right, rather than their undesirability being contingent on any downstream effect. The following imaginative exercise, I suggest, further supports that view. Consider your ideal conception of a socially just society. Is it a matter of indifference to you whether (i) the members of that society accord default parity of esteem between their fellows, or (ii) the members of that society possess attitudes that rank and grade their fellows prior to encountering them, but suppress these attitudes in order to act in an outwardly even-handed way? Clearly, there is something about the second scenario that offends our conception of justice. The attitudes characteristic of social equality are important in virtue of being necessary for a socially just society. Profiling beliefs, in so far as they preclude the realisation of social justice, are of intrinsic disvalue.

To close, it is worth acknowledging that social equality would sometimes be served by lacking access to various types of adverse demographic evidence. This alone might suffice to undermine recent arguments in favour of demographic differences research.²¹ However, the gathering and disseminating of adverse demographic evidence can serve a legitimate socio-political function. Democratic accountability is one example. The fact that some negative characteristic—for example, alcoholism or criminality—is statistically prevalent within some demographic can underscore valuable criticisms of how certain groups are treated, identify public spending priorities, and measure progress between strategies for addressing problems. These critical uses of demographic evidence can ultimately be productive of social equality. Given that such evidence can serve a legitimate function, it is incumbent on us to provide a precise characterisation of how we should respond to it in our daily lives.

5. Profiling and Epistemic Neutrality

I now outline a positive response to the problem of profiling. The ambition is to vindicate a negative assessment of profiling within a traditional conception of moral and epistemic normativity—one denying the susceptibility of belief to moral evaluation, appealing only to truth-related factors in determining epistemic rationality. To develop this account, I will draw on the attitudes of *epistemic neutrality*.

When discussing social equality, we reflected upon the importance of not harbouring antecedent views about whether an individual is more, or less, likely to have certain virtues or vices before having evidence pertaining to their individual conduct. One way to avoid drawing pernicious distinctions is to suppose that is simply *equally likely* for any individual to possess the relevant characteristics. But this response—which might require believing *against* the evidence—is not the only way for attitudes to conform with social equality. In addition to endorsing or rejecting a proposition, there is a third type of response open to us. These are the epistemically neutral attitudes.

Neutrality is different from the mere absence of any cognitive attitude. Inanimate objects lack cognitive attitudes towards every proposition without being epistemically neutral. Neutrality, as the Pyrrhonians knew well, is a *positive* state of mind. Despite

²¹ For defence of demographic differences research, see Carl [2018] and Cofnas [2020].

the familiarity of these neutral states, only recently has the nature and normative profile of neutrality been subjected to rigorous analysis (see Friedman [2013, 2017, 2019], Rosa [2021], and McGrath [forthcoming]).

McGrath [ibid.] helpfully distinguishes between three different types of epistemic neutrality—agnosticism, suspension, and inquiry:

Agnosticism is having an intermediate confidence about some proposition.

Suspension is intentionally omitting judgment on some proposition, whether with a view to judging later, or to be refraining from judging.

Inquiring is the goal-directed state of mind that we are in when trying to answer a question.

An important *desideratum* for my account is that our response to profiling should not involve believing (or holding a credence) against the evidence. Agnosticism can naturally be represented as a middling credence that p , or as a full belief that p is (roughly) just as likely as not. If there can be profiling cases involving strong evidence, agnosticism is not a promising way to stay within the bounds of purely truth-related epistemic rationality. However, the normative profile of suspension and inquiry are more distinctive. It is to these, particularly to suspension of judgment, that we should turn in responding to the problem of profiling.

Suspension of judgment is interestingly unlike other epistemic states—like full beliefs or degrees of confidence—because it involves a commitment *not* to endorse any particular conclusion about some question. Suspension of judgment can serve a variety of different roles. We might suspend judgment in order to continue to inquire, to postpone inquiry, or to terminate inquiry indefinitely. A lesson from recent work is that suspension is *normatively* distinctive. While beliefs are typically assessed against the strength of your evidence at a particular time t , suspension can seemingly be justified by pragmatic factors. For example, you might be deciding which car to purchase, and the evidence at t might seem to strongly favour car A over car B. But suppose that you asked a car-savvy friend for advice, and you see that she has just sent you a long email about the merits of the two models. It seems justifiable to suspend judgment on which car to buy, until you have energy to read the email at a later time. This reason to suspend judgment concerns, not how strong your evidence is at t , but rather the pragmatic fact that better evidence is available in the future.²² Notably, there is a sense in which you are justified in suspending judgment *even if* the judgment about which car to buy, once formed, would have been epistemically justified at t , given the fact that the evidence strongly favoured car A over car B. Such cases suggest that the permissibility of suspending judgment is, in fairly routine cases, not undermined by the epistemic permissibility of forming a judgment on the question under discussion. For example, while we might have strong evidence that p at t , the prospect of having utterly decisive evidence whether p in an hour's time seems to be a good reason to suspend judgment, notwithstanding the epistemic permissibility of believing p at t . The amenability of suspended judgment to non-evidential justification is further supported by the existence of pragmatic reasons *not* to suspend. For instance, if matters are urgent, suspending judgment might be rendered

²² There are other examples, such as where you are temporarily impaired at t : again, it seems justified to suspend judgment until the impairment passes. For discussion, see McGrath [forthcoming]. Additional cases are provided by Schroeder [2012].

irrational by countervailing pragmatic reasons: you must decide *now*, based on the evidence currently available.

I claim that not only does suspended judgment admit of pragmatic justification, but, crucially, that it also admits of *moral* justification. I will illustrate this with a case.

Panel. You are browsing the dossier of a candidate whom you will, later that day, interview for an academic position. On reading the candidate's name—James Smith—you have a nagging feeling that a philosopher with this name was embroiled in a torrid affair with a friend of yours. But you can't be sure. You had *only* heard about this disreputable romance from someone else, via email, and the name was inessential to the conversation. Doubts enter. James Smith is a common name. You might be mixing it up with *Joe* Smith, or with *another* James Smith. This would easily be settled by opening your email. You know that if you do make up your mind that it *was* him, then you'll be unfairly hostile. You take your ethical responsibilities as interviewer seriously, and so, despite your curiosity, you decide that it would be best to suspend judgment. You dwell on your doubts, and then commit to not thinking about it until afterwards, distracting yourself now by opening their writing sample.

Suppose that you cannot recuse yourself at this late stage. In such a case, you have a moral reason to suspend judgment on Smith's involvement in the disreputable romance. The moral consideration—namely, fairness to a job candidate—is a distinctive feature of the case. It would not be present, for instance, if you had no interviewing duties; then you would be free to make up your mind on the basis of recollection, or to check the email sent by your friend.

By focusing on moral reasons favouring suspending judgment, we capture the important data-point that an agent can be open to criticism in virtue of *forming* a particular belief—here, because they ignore the moral reasons in favour of suspending judgment. Holding that suspended judgment admits of moral justification is entirely consistent with its being the case that, once formed, your belief would *only be assessable against the strength of the evidence supporting it*. The choice of whether or not to suspend judgment is a distinct *locus* of moral assessment, separate from supposing that the epistemic rationality of belief is encroached upon by associated moral costs, or that belief itself is a subject of moral evaluation.

Suspension of judgment, I suggest, is the appropriate response to cases in which we have only demographic statistics speaking to the likelihood of someone possessing some particular virtue, vice, or level of intelligence. Return to the *Diner* case. My claim is that the correct response is to suspend judgment about whether any particular diner is responsible, or even whether they are more likely to be responsible, for the theft. Suspension of judgment in the *Diner* case can either be permanent or temporary, depending on whether further non-statistical evidence will become available in the future.

The moral reasons in favour of suspension here are (at least) two-fold: there are the moral risks of exposing members of a certain demographic to undeserved heightened scrutiny, and, second, since relational equality is an important value that well-intentioned agents should be concerned with promoting, there is reason to avoid forming attitudes that are destructive of relational equality. These are forceful moral reasons to suspend judgment. They obtain, irrespectively of whether the profiling belief, if formed, would be epistemically rational.²³

²³ The same applies, even if the profiling belief would count as knowledge. For discussion of probabilistic knowledge and profiling, see Moss [2018].

Although focusing on individual cases can make it sound as if the decision to suspend must be taken anew, every time that we face profiling scenarios, a better view is that we ought to adopt a *default policy* of suspending judgment when all that we have to go on is adverse demographic information pertaining to virtue, vice, or intelligence. As has been effectively explained by Bratman [1989], we often constrain deliberation in advance by committing to φ (for example, *wear a seatbelt*) without weighing up costs and benefits every time that we face a φ -relevant decision. Adopting such a pre-commitment to suspend judgment in profiling scenarios exhibits a standing concern for social equality. It also provides a satisfying perspective on an issue noted earlier with the risk-based moral encroachment view: namely, it seems somewhat pathological to weigh the costs and benefits of profiling every time that we encounter such cases.

So, the view is that a default policy of suspending judgment is the best response to cases in which demographic statistics could support an inference as to the intelligence, virtues, or vices of a member of that demographic. This includes, under my account, both positive and negative profiles. I find it hard to conceive of legitimate exceptions to this policy with respect to these characteristics, even if it is fairly straightforward to find examples where *other* sorts of profiling are justified. For instance, it seems permissible for a medical practitioner to use statistical information to profile someone for a demographically prevalent illness such as sickle cell disease. Whether or not there are such exceptions for the view that I have outlined here is a challenge that I leave for opponents to pursue.

By way of considering three important objections, I want to end by reflecting on some broader normative issues raised by my suspension-based solution to the problem of profiling.

First, since I describe judging and failing to suspend in broadly agential terms, doxastic voluntarism worries re-emerge. We should not overstate such worries. Clearly, we *do* suspend judgment on some questions. Regardless of whether we have direct control over our judgments, we certainly have various types of indirect control. There are many practical ways in which we can prime ourselves to suspend: for instance, if we want to avoid forming a judgment about the misdeeds of a friend, we might interrupt a story that looks like it will paint them in a negative light. There are cognitive elements to our control, too. For example, Weatherson [2008] effectively shows that we have control over our stock of beliefs, since we are able to deliberately identify doubts and bring ourselves to a point of uncertainty about propositions that we are evaluating. These aspects of control are evident in the *Panel* case: we can bring ourselves to suspend judgment by intentionally omitting to gather further evidence, and by consciously entertaining doubt. In profiling cases, it is generally easy to identify and entertain doubts that are apt to prime suspension of judgment. This is because any individual might easily be an exception to a statistical generalisation about the group of which they are a member. This doubt is straightforward to identify and to hold in mind. Humans vary in significant and unpredictable ways, even if some adverse characteristic is predominant within a group. Indeed, the extreme variability of humans is an important difference between us and other categorizable entities, such as tokens of a particular type of car.²⁴

²⁴ See Munton [2019] for discussion.

A second worry is derived from the apparent peculiarity in my view that we can morally criticise agents for failing to suspend, but we can deny that beliefs born from this failure to suspend are open to moral evaluation. Indeed, given that I argue that profiling beliefs are destructive of social equality, one might wonder how we can make sense of the exclusion of belief from moral evaluation.

In response, I suggest that the moral evaluation of belief is a category error. Agents can be rightly criticised for *forming* certain beliefs, even if beliefs themselves are not a proper *locus* of moral evaluation. The fact that the creation and maintenance of something is assessable against a norm *n* does not mean that the thing itself is assessable against *n*. A simple example illustrates this. Clearly, it would be morally better, had the Falkland Islands not been strewn with landmines.²⁵ Everyone can accept that. But this doesn't mean that the landmines themselves were a *locus* of moral evaluation, even if their presence was morally detrimental. (Landmines have the function of exploding when stepped upon. The norms of being a good landmine are derivative of this function.) Rather, we morally evaluate the people who put the landmines there and/or the people who failed to remove them. Even if we were to felicitously call the Falkland landmines 'morally bad', this would be an attenuated way of saying that it is bad that they were placed, that they were not removed, or that they are liable to have certain bad effects.²⁶ It would be a category error to call the landmine itself a subject of moral evaluation. I think that something similar holds for the evaluation of belief. The aim of belief is truth (or knowledge): the norms of epistemic rationality are derivative of that aim. What *is* morally evaluable is your choosing to manage (namely, create and maintain) your stock of beliefs in a way that evinces disregard for (i) the importance of social equality and (ii) downstream risky effects of holding profiling beliefs. This preserves the idea that moral evaluation pertains only to our actions, omissions, and character. Allowing oneself to form a judgment can be a type of action—or, at least, failing to bring oneself to suspend can be a type of omission—and such actions and omissions are a proper *locus* of moral evaluation.

Finally, I want to reflect upon and forestall misunderstanding about where this paper leaves us with respect to the distinction between moral and epistemic normativity.

An initial worry is a narrow one as to whether suspension might violate some epistemic norm prohibiting believing against the evidence. My answer to this worry is that suspension of judgment is *not a type of belief*; it is one's intentionally prescinding from believing, a different type of cognitive act. So, even if the evidence in favour of the candidate belief is strong, we don't believe contrary to the evidence if we suspend. However, recommending suspension in profiling cases does not require us to reject a normative or prescriptive reading of the idea that our beliefs should aim at truth or knowledge. In so far as it is popular to endorse *conditional* readings of such norms—for example, that the (apparent) truth of a proposition *p* is a necessary but not sufficient condition for being obligated to believe *p*—we can retain normative readings of the aim of belief while still leaving room to permit suspension of judgment.

Although my account is traditional by holding that belief is amenable only to epistemic assessment, it accommodates the thought that non-epistemic matters impinge

²⁵ Historical note: while this paper was under review, the Falklands were in fact finally declared landmine-free.

²⁶ This shorthand appears in other evaluative domains: e.g. we say that a textbook 'contains knowledge' even though the textbook is not an epistemic subject.

upon the proper conduct of our doxastic lives, more broadly. For example, I argue that the decision *not* to form a given belief *can* be legitimised by moral factors. So, in this sense, moral reasons favouring suspension can weigh on the opposite side of the scale against epistemic reasons to form a belief. However, the existence of these moral reasons bearing on how we regulate our intellectual lives is in fact entirely unsurprising, because the relevance of non-epistemic factors in such matters is routine and prevalent. This can be seen clearly with inquiry, another of the neutral attitudes. Sometimes we are justified to stop inquiring into some issue because, for example, it is bad for our health (pragmatic), or because the fruits of our inquiry will be misapplied and put to malign ends (moral). Things are much the same with suspension: it is amenable to moral and pragmatic justification. Thus, moral and pragmatic factors matter for the broader question of how we conduct ourselves as thinking agents, even though it would be a category error to say that beliefs born from a failure to suspend are themselves a *locus* of moral evaluation.

6. Conclusion

Traditional views that preserve a sharp distinction between epistemic evaluation of belief and moral assessment of action and character are compatible with rejecting demographic profiling. The moral failures involved in demographic profiling can be located in the decision not to suspend judgment, rather than supposing that beliefs themselves are a *locus* of moral evaluation. A key moral reason to suspend judgment when faced with adverse demographic evidence is to promote social equality—explaining why positive profiling is dubious, in addition to familiar cases of negative profiling, and why profiling is suspect even when no particular action is at stake. This view, although compatible with revisionary normative views, does not presuppose them. Hence, philosophers of all stripes can reject demographic profiling both in thought and in deed.²⁷

Disclosure Statement

No potential conflict of interest was reported by the author.

ORCID

Lewis Ross  <http://orcid.org/0000-0001-9683-9924>

References

- Anderson, Elizabeth S. 1999. What Is the Point of Equality? *Ethics* 109/2: 287–337.
- Backes, Marvin 2020. Epistemology and the Law: Why There Is No Epistemic Mileage in Legal Cases, *Philosophical Studies* 177: 2759–78.
- Basu, Rima 2019a. The Wrongs of Racist Beliefs, *Philosophical Studies* 176/9: 2497–515.
- Basu, Rima 2019b. Radical Moral Encroachment: The Moral Stakes of Racist Beliefs, *Philosophical Issues* 29/1: 9–23.

²⁷ I would like to thank, for generous comments on earlier drafts, Tuuli Ahlholm, Claire Field, Matt McGrath, Martin Smith, Alex Voorhoeve, anonymous referees, the Editor and an Associate Editor of this journal. Special thanks to Matt McGrath for conversations, spanning two continents, instrumental in developing this paper.

- Basu, Rima and Mark Schroeder 2019. Doxastic Wronging, in *Pragmatic Encroachment in Epistemology*, ed. Brian Kim and Matthew McGrath, New York: Routledge: 181–205.
- Bolinger, Renee Jorgensen 2020a. The Rational Impermissibility of Accepting Racial Generalizations, *Synthese* 197/6: 2415–31.
- Bolinger, Renee Jorgensen 2020b. Varieties of Moral Encroachment, *Philosophical Perspectives* 34/1: 5–26.
- Bratman, Michael E. 1989. Intention and Personal Policies, *Philosophical Perspectives* 3: 443–69.
- Buchak, Lara 2014. Belief, Credence, and Norms, *Philosophical Studies*, 169/2: 285–311.
- Carl, Noah 2018. How Stifling Debate Around Race, Genes and IQ Can Do Harm, *Evolutionary Psychological Science* 4/4: 399–407.
- Chan, Timothy, ed. 2013. *The Aim of Belief*, Oxford: Oxford University Press.
- Clifford, William, 1877 (1974). The Ethics of Belief in *The Ethics of Belief and Other Essays*, ed. Timothy Madigan, Amherst, MA: Prometheus: 70–96.
- Cofnas, Nathan 2020. Research on Group Differences in Intelligence: A Defense of Free Inquiry, *Philosophical Psychology* 33/1: 125–47.
- Cohen, Gerald A. 1989. On the Currency of Egalitarian Justice, *Ethics* 99/4: 906–44.
- DeRose, Keith 1992. Contextualism and Knowledge Attributions, *Philosophy and Phenomenological Research* 52/4: 913–29.
- Di Bello, Marcello and Collin O’Neil 2020. Profile Evidence, Fairness, and the Risks of Mistaken Convictions, *Ethics* 130/2: 147–78.
- Ebert, Philip A., Martin Smith, and Ian Durbach 2018. Lottery Judgments: A Philosophical and Experimental Study, *Philosophical Psychology* 31/1: 110–38.
- Enoch, David, Levi Spectre, and Talia Fisher 2012. Statistical Evidence, Sensitivity, and the Legal Value of Knowledge, *Philosophy & Public Affairs* 40/3: 197–224.
- Fourie, Carina, Fabian Schuppert, and Ivo Wallimann-Helmer, eds, 2015. *Social Equality: On What It Means to Be Equals*, New York: Oxford University Press.
- Friedman, Jane 2013. Suspended Judgment, *Philosophical Studies* 162/2: 165–81.
- Friedman, Jane 2017. Why Suspend Judging? *Noûs* 51/2: 302–26.
- Friedman, Jane 2019. Inquiry and Belief, *Noûs* 53/2: 296–315.
- Fritz, James 2017. Pragmatic Encroachment and Moral Encroachment, *Pacific Philosophical Quarterly* 98/1: 643–61.
- Hawthorne, John 2004. *Knowledge and Lotteries*, Oxford: Clarendon Press.
- Hájek, Alan 2007. The Reference Class Problem Is Your Problem Too, *Synthese* 156/3: 563–585.
- Gardiner, Georgi 2018. Evidentialism and Moral Encroachment, in *Believing in Accordance with the Evidence*, ed. Kevin McCain, Cham: Springer Verlag: 169–95.
- Gardiner, Georgi 2019. The Reasonable and the Relevant: Legal Standards of Proof, *Philosophy & Public Affairs* 47/3: 288–318.
- Garrau, Marie, and Cecile Laborde 2015. Relational Equality, Non-Domination, and Vulnerability in *Social Equality: On What It Means to Be Equals*, ed. Carina Fourie, Fabian Schuppert, and Ivo Wallimann-Helmer, New York: Oxford University Press: 45–64.
- Jackson, Elizabeth and James Fritz, [forthcoming](#). Belief, Credence, and Moral Encroachment, *Synthese*.
- Leslie, Sarah-Jane 2017. The Original Sin of Cognition: Fear, Prejudice, and Generalization, *The Journal of Philosophy* 114/8: 393–421.
- Lippert-Rasmussen, Kasper 2018. *Relational Egalitarianism: Living as Equals*, Cambridge: Cambridge University Press.
- Littlejohn, Clayton 2020. Truth, Knowledge, and the Standard of Proof in Criminal Law, *Synthese* 197: 5253–86.
- Liu, Xiaofei and Liang Ye [forthcoming](#). What It Means to Respect Individuality, *Philosophical Studies*.
- McGrath, Matthew [forthcoming](#). Being Neutral: Suspension of Judgment, Agnosticism and Inquiry, *Noûs*.
- Miller, David 1997. Equality and Justice, *Ratio* 10/3: 222–37.
- Mogensen, Andreas 2019. Racial Profiling and Cumulative Injustice, *Philosophy and Phenomenological Research*, 98/2: 452–77.
- Moss, Sarah 2018. Moral Encroachment, *Proceedings of the Aristotelian Society* 118/2: 177–205.
- Munton, Jessie 2019. Beyond Accuracy: Epistemic Flaws with Statistical Generalizations, *Philosophical Issues* 29/1: 228–40.

- Nunn, G. Alexander 2015. The Incompatibility of Due Process and Naked Statistical Evidence, *Vanderbilt Law Review* 68: 1407–33.
- Pace, Michael 2011. The Epistemic Value of Moral Considerations: Justification, Moral Encroachment, and James' 'Will To Believe', *Noûs* 45/2: 239–68.
- Pritchard, Duncan 2018. Legal Risk, Legal Evidence and the Arithmetic of Criminal Justice, *Jurisprudence* 9/1: 108–19.
- Pundik, Amit 2008. Statistical Evidence and Individual litigants: A Reconsideration of Wasserman's Argument from Autonomy, *The International Journal of Evidence & Proof*, 12/4: 303–24.
- Redmayne, Michael 2008. Exploring the Proof Paradoxes, *Legal Theory* 14/4: 281–309.
- Reichenbach, Hans 1949. *The Theory of Probability*, Berkeley: University of California Press.
- Rosa, Luis 2021. Rational Requirements for Suspended Judgment, *Philosophical Studies* 178/2: 385–406.
- Ross, Lewis D. 2020. Recent Work on the Proof Paradox, *Philosophy Compass* 15/6: e12667.
- Ross, Lewis 2021. Rehabilitating Statistical Evidence, *Philosophy and Phenomenological Research* 102/1: 3–23.
- Scheffler, Samuel 2015. The Practice of Equality in *Social Equality: On What It Means to Be Equals*, ed. Carina Fourie, Fabian Schuppert, and Ivo Wallimann-Helmer, New York: Oxford University Press: 21–44.
- Schroeder, Mark 2012. The Ubiquity of State-Given Reasons, *Ethics*, 122/3: 457–88.
- Smith, Martin 2018. When Does Evidence Suffice for Conviction? *Mind* 127/508: 1193–218.
- Wasserman, David T. 1992. The Morality of Statistical Proof and the Risk of Mistaken Liability, *Cardozo Law Review* 13: 935–76.
- Weatherston, Brian 2008. Deontology and Descartes's Demon, *The Journal of Philosophy* 105/9: 540–69.
- Wolff, Jonathan 1998. Fairness, Respect, and the Egalitarian Ethos, *Philosophy & Public Affairs* 27/2: 97–122.
- Wolff, Jonathan 2019. Equality and Hierarchy, *Proceedings of the Aristotelian Society* 119/1: 1–23.
- Young, Iris Marion 1990. *Justice and the Politics of Difference*, Princeton: Princeton University Press.