

RACIAL PROFILING AND CUMULATIVE INJUSTICE¹

Andreas Mogensen

University of Oxford

1. Introduction

Many people believe that racial profiling is wrong. A 2004 Gallup poll of public attitudes on racial profiling found that most Americans regard it as unjustified for each context about which respondents were surveyed. For example, 67% expressed the opinion that police may never treat race as a predictor of crime when deciding whether to stop motorists on roads or highways.²

What exactly is it that makes racial profiling of this kind wrong? I believe that philosophers working in applied ethics and in legal and political philosophy have yet to work out a fully satisfactory answer to this question. Obviously, some are not trying to: one of the most cited articles in the literature is a paper by Mathias Risse and Richard Zeckhauser (2004) arguing that profiling is permissible under certain real-world conditions.³ But even philosophers who argue against Risse and Zeckhauser seem to fail in

¹ I'm very grateful to Ben Eidelson and Tom Sinclair for their comments on previous drafts of this paper, and to an anonymous referee at this journal for their helpful criticisms and suggestions.

² Carlson (2004). 72% deemed racial profiling unjustified as a tactic for apprehending shoplifters and 53% deemed it unjustified as an airport security measure. For more on public attitudes to racial profiling, see Boonin (2011: 300-1).

³ See also Boonin (2011), Levin (1992). Central to Risse and Zeckhauser's argument is the suggestion that harm-based arguments of the kind to be discussed in section 3 misattribute to racial profiling harms actually

capturing the sense of pre-theoretic opposition that so many of us feel when it comes to racial profiling: as we'll see, these philosophers have had some difficulty in explaining what might be wrong with racial profiling that occurs without public knowledge.

In this paper, I'll consider how we might provide a more satisfactory explanation for why racial profiling involves a serious injustice. Like many working in this area, I will focus on racial profiling of black Americans by US law enforcement. One important part of the argument to be developed in this paper will be to insist that racial profiling is *pro tanto* wrong in and of itself and violates a constraint on fair treatment, a constraint that's violated quite generally by acts of statistical discrimination that rely on ascribed characteristics. But violation of this constraint isn't of itself an especially serious wrong and can easily be counterbalanced by opposing reasons. To fully capture the significant wrong that occurs when racial profiling is targeted at black Americans or other similarly situated groups, we should think of this basic injustice as exacerbated when it can be viewed as part of a larger pattern of similar actions that collectively bring about or realize a very bad or morally catastrophic outcome. It will be argued that these pattern-based reasons⁴ are key in fully explaining and characterising the serious injustice involved in racial profiling of black Americans and other similarly situated groups, drawing on ideas worked out by theorists of oppression influenced by Marilyn Frye's (1983) birdcage analogy.⁵

Here is the plan. Section 2 will clarify how I mean to use the term 'racial profiling'. Section 3 will identify the key challenges for a satisfactory philosophical account of the

caused by structural racism, a mistake occasioned by the fact the former may be perceived as a salient symbol of the latter. For reasons like those noted by Lever (2005) and Eidelson (2015: 202-6), I find this misattribution hypothesis implausible.

⁴ I borrow the term 'pattern-based reasons' from Woodard (2008).

⁵ See Lippert-Rasmussen (2014: 170) and Brennan (2016: 246-7) for similar ideas in the recent discrimination literature.

wrong in racial profiling. Section 4 will develop a positive proposal to the effect that racial profiling is intrinsically unjust. Section 5 will build on this idea and work out the key connections to relevant philosophical work on collective wrongdoing and oppression. Section 6 will answer some potential objections. Section 7 will offer a summary and conclusion.

2. Some points of clarification

For the purposes of this paper, I'm going to define 'racial profiling' such that law enforcement officers engage in racial profiling if and only if they subject individuals belonging to certain ethnic or racial categories to increased scrutiny, intrusion, *etc.* based on the expectation (explicit or implicit) that members of such groups are statistically overrepresented among those involved in certain types of criminal behaviour. Note that it's no part of this definition that profiling occurs only if race is the *sole* basis for suspicion: racial profiling may use race in conjunction with other descriptors.^{6 7}

As is standard, I don't count all cases where race serves as a basis for increased suspicion as instances of racial profiling. In particular, the above definition does not cover the use of race or ethnicity as a basis for increased suspicion when police are trying to

⁶ According to Meehan and Ponder (2002), racial profiling "is typically defined as the police use of race as the *sole* basis for initiating law enforcement activity." (403) Harris (2003) often writes in a way that suggests he thinks of racial profiling as an alternative – and not a complement – to behavioural profiling.

⁷ Additionally, it is worth emphasizing that, by this definition, the expectation that members of some racial group are overrepresented amongst those involved in some type of criminality need not rest on the odious supposition that a person's race predisposes them toward criminal behaviour. For example, someone holding this expectation may assume that its explanation rests with the fact that historical and contemporary systems of privilege and oppression result in members of certain racial groups becoming overrepresented among those experiencing negative socioeconomic conditions under which the incentives to engage in certain types of criminality are greater and the disincentives weaker. On this point, see Boonin (2011: 309).

apprehend someone matching a suspect description that includes information about the race or ethnicity of the perpetrator.⁸ This sort of activity doesn't constitute racial profiling under the proposed definition, as it doesn't require any assumptions about whether members of the relevant demographic category are statistically overrepresented among those involved in certain types of criminal behaviour. In particular, it's quite compatible with denying that assumption. Police searching for a mugger identified by the victim as South Asian might also believe that South Asian people are less likely as a rule to be involved in this sort of crime.⁹

As it is and has been practiced, racial profiling is associated with a range of serious injustices. The expectation that members of certain racial or ethnic groups are statistically overrepresented among those involved in certain types of criminal behaviour may reflect nothing more than racial prejudice.¹⁰ Insofar as that expectation is based on statistics, these may be of questionable evidential import.¹¹ Racial profiling may be used to enforce unjust laws or in conjunction with 'high-discretion' police tactics that provide officers with excessive latitude in their ability to detain and harass ordinary citizens.¹² Racial profiling often involves abusive behaviour that is seriously wrong quite apart from any wrong inherent in racial profiling: in the worst instances, racially biased police interventions result

⁸ In this respect, my definition agrees with that proposed by Ramirez et al. (2000). For more on the distinction between the use of race in profiling and suspect descriptions, see Eidelson (2015: 177-187).

⁹ I base this example on Eidelson (2015: 183).

¹⁰ But see Harris (2003: 100-102), Kennedy (1999) for valuable correctives to the idea that racial profiling is simply a matter of individual prejudice.

¹¹ Consider, for example, the apparent circularity of using racial differences in arrest and conviction rates that may simply be artefacts of racial profiling to justify the continuation of profiling. On this issue, see Boonin (2011: 313-4), Harris (2003: 75-8), Zack (2015: 46-8).

¹² See Harris (2003: 30-47) on the interrelationship between racial profiling and high-discretionary police tactics.

in officers wrongfully killing black Americans. To make matters worse, the existing legal system appears unable to effectively address the abuse and violence suffered by black Americans at the hands of law enforcement, as evinced, for example, by the dismaying failure of recent grand juries to indict the police officers who killed Michael Brown, Tamir Rice, and Eric Garner.¹³

The injustices described in the previous paragraph are entirely compatible with racial profiling as I have defined it, but they are not implied by that definition. It is common in the philosophical literature to try to abstract away from the kind of wrong-making features I have listed, the aim being to ask questions that pertain to the morality of racial profiling considered ‘on its own terms’, as opposed to as the locus for obvious injustices of this sort. Following Kasper Lippert-Rasmussen (2006), we can define *unalloyed racial profiling* as racial profiling that reflects accurate and well-evincing statistical regularities, deployed in an unprejudiced way via otherwise acceptable police tactics that enforce justified legal requirements, serving thereby to reduce the incidence of crime. (Note, however, that it is no part of the definition of unalloyed profiling that it does not occur in a society that has been and continues to be characterized by systemic racism.) The philosophical focus tends to be on unalloyed profiling, so understood.¹⁴ The question is whether profiling of this kind would be wrong or whether the wrongness of racial profiling as practiced is exhaustively explained in terms of the obvious injustices noted in the previous paragraph. Public opposition to racial profiling plausibly cannot be exhaustively explained in these terms.¹⁵ Those who oppose racial profiling of black Americans are generally not in favour of simply

¹³ See Zack (2015: 63-91).

¹⁴ For this way of framing the philosophical debate see, *inter alia*, Boonin (2011: 301-7), Hellman (2014), Lippert-Rasmussen (2006), Risse & Zeckhauser (2004).

¹⁵ Boonin (2011: 303-7).

reforming the way it is practised or limiting the conditions under which it may be used:¹⁶ they are in favour of abolishing it outright.¹⁷ Public opposition to racial profiling plausibly carries through to the very idea of rendering black Americans subject to increased suspicion on account of their race.

Here is one final point of clarification. It is no part of my definition that racial profiling has to be the explicit, official policy of those law enforcement agencies in which it is practised. This assumption is often adopted in debates on racial profiling and key objections that have been raised against racial profiling seem to rely on this assumption, as we'll see. This leaves us with a significant blind spot: namely, the possibility of racial profiling that is not acknowledged policy and occurs without public knowledge.

There is nothing altogether extraordinary about criminal profiling that relies on descriptors withheld from public knowledge. The Computer Assisted Passenger Pre-Screening (CAPPS) program introduced in the 1990s to screen for terrorists at US airports relied on a computerized profile most details of which were never released to the public, though the Department of Justice insisted that race did not feature on the list.¹⁸ As Benjamin Eidelson (2015) notes, “in most contexts ... opaqueness about just what triggers increased scrutiny is the norm – indeed, opaqueness is essential to the utility of the

¹⁶ The 2004 Gallup survey on racial profiling asked if racial profiling is “*ever* justified” (my emphasis). A majority answered “No, is *never* justified” (again, my emphasis) for all three contexts about which they were asked.

¹⁷ This may be explained, however, by the presumption that adequate reform is simply impossible given the present state of things.

¹⁸ Harris (2003: 139–144).

profiling.” (190) The possibility that racial profiling might conceivably be carried out in secret is worth keeping in mind, therefore.¹⁹

3. Why is racial profiling wrong?

Many of us have a pre-theoretic sense that there is something wrong *in principle* with racial profiling. We are likely to agree with George C. Thomas (2001) that “[s]ingling people out on the basis of race, is a violation of the right to equal concern and respect” (43) and with Randall Kennedy’s (1999) insistence that “individuals should be judged by public authority on the basis of their own conduct and not on the basis—not even *partly* on the basis—of racial generalization.” (34) Those of us who are so disposed regard racial profiling as intrinsically unjust: we think of it as *pro tanto* wrong even apart from its capacity to generate such characteristic harmful effects as the stigmatization of profiled groups or a breakdown in trust between minority communities and police. We think it would be wrong to engage in racial profiling even if no one knew that racial profiling was being practiced, whereas this would seem to obviate harmful effects like those just mentioned. And we think this even with respect to contexts where there might otherwise be a reasonable case for criminal profiling in secret.

A view of this sort is subject to significant philosophical challenges, however. In particular, it is difficult to maintain in a way that avoids overgeneralizing, leading us to condemn practices that we think of as involving no significant wrongdoing.

To see why, we can begin by noting that very few of us are opposed to criminal profiling *per se*. Although we are incensed by racial profiling, we don’t feel a similar sense of injustice at the idea that airport security officials monitoring for potential drug couriers

¹⁹ For more on covert racial profiling as a practical possibility, see Boonin (2011: 338–9) and Eidelson (2015: 215–9).

might give increased scrutiny to passengers who pay in cash, check no hold luggage, and make frequent short-stay trips to key drug source cities. Plausibly, we regard race in particular as an inappropriate basis for criminal profiling.²⁰ What marks out race as especially inappropriate in this way?

It is tempting to answer that race is different because race is unchosen and thereby constitutes a morally unacceptable basis for statistical discrimination. Judging people on the basis of things which lie beyond their control may seem uniquely unfair. However, this reply seems less than fully convincing once we note that we aren't generally opposed to statistical discrimination that relies on other similarly unchosen traits.

Consider the following examples, all noted previously in the philosophical literature. Many of us are okay with the idea that young men may be charged more for car insurance in light of the fact that they show a greater statistical propensity for risky driving.²¹ We also accept that age can be used as a crude indicator in discriminating between those citizens who are competent to enjoy the right to vote and those who aren't.²² And we accept that age can be used in a similar way in picking out the class of people deemed sufficiently responsible to enjoy the right to drive (though here we also apply a separate test of driving competence).²³ In all of these cases, we allow individuals to be judged and disadvantaged on the basis of generalizations about their unchosen traits.

Obviously, none of the foregoing examples involve race, but there are other cases that suggest we may not be strongly opposed in principle to statistical discrimination that relies on race. Certain affirmative action programs may be thought to involve disadvantaging people on the basis of statistical generalizations about their race or

²⁰ Eidelson (2015: 188-9).

²¹ For this point, see Lippert-Rasmussen (2007: 385), Wasserman (1991: 946).

²² On voting and statistical discrimination by age, see Schauer (2003: 118-121).

²³ Lippert Rasmussen (2007: 395).

ethnicity, for example.²⁴ Consider a university admissions program similar to that used by the University of Texas at Austin that allows assessors to give preference to students in light of their racial or ethnic background when doing so may be expected to serve relevant educational goals such as breaking down stereotypes and promoting cross-racial understanding. We might suppose that a minority candidate's race is a reliable statistical indicator that her admission will serve these goals better than an otherwise similar white candidate, but not an absolute guarantee.²⁵ If that's correct, the admitting institution seems to treat race as a statistical predictor that some students are better placed than others to contribute to the institution's educational goals.

Of course, this case isn't going to elicit the same level of agreement as the examples involving auto insurance, voting, or the right to drive. Although UT Austin's program has been upheld by the Supreme Court, affirmative action continues to be hugely controversial and many reject it precisely because they feel it violates a general principle of colour-

²⁴ On affirmative action and racial profiling, see Kennedy (2002), Levin (1992). See also Goldman (1979: 100-101).

²⁵ This may depend on what it means to, say, 'break down stereotypes' in the intended sense. We could interpret this phrase in such a way that breaking down stereotypes is something that can only be done by members of the stereotyped group insofar as they frustrate stereotyped expectations about their own attitudes and behaviour. In that case, the stated supposition might seem inappropriate, since a candidate who does not belong to the stereotyped group simply could not contribute to this goal. I am assuming a broader understanding of what it means to 'break down stereotypes' where this is something to which others besides those who are stereotyped could contribute: e.g., by calling out stereotyped expectations that members of other groups might air only when they know they won't be overheard by those whom they stereotype. Regardless of how this interpretative question might be settled in any given case, my expectation is that there are few whose views about the permissibility of affirmative action turn decisively on whether it is settled in one way as opposed to the other. I'm grateful to an anonymous referee for pressing me on this issue.

blindness.²⁶ Even so, many who feel that there is something wrong in principle with racial profiling also stand behind affirmative action programs: the present author, for example. What's more, a blanket requirement of colour-blindness is simply implausible. To vary a stock example, a casting director can obviously search for a black actor to play Ruth Younger in a production of *A Raisin in the Sun*. She can also permissibly search for a white actor to play Karl Lindner in the same production.²⁷ Nor do we insist on total colour-blindness in policing, as many see nothing inherently wrong in allowing police to use race as an identifying factor when relying on a suspect description offered by a witness.²⁸ It seems we don't in fact believe that singling out people on the basis of their race necessarily violates a right to equal respect. Instead, our intuitions are sensitive to the manner in which race comes to figure as a basis for increased suspicion.

Hitting even closer to home, it's notable that few²⁹ have detected any significant injustice in the use of profiles that identify yet-to-be-caught serial killers as white, as happened during the manhunt for the Beltway sniper in 2002.³⁰ It seems that we can't really count ourselves as opposed to racial profiling *per se*, but only to the profiling of certain groups in certain contexts. And this can seem hard to square with the view that racial profiling is wrong because it is intrinsically unjust.

By contrast, the kind of selectivity that we've just observed looks easy to explain if we think that racial profiling targeted at black Americans or other similar minority communities is wrong because it engenders harmful effects like increased stigmatisation.

²⁶ See Kennedy (2013) for an overview of recent controversy on affirmative action.

²⁷ Cf. Singer (1978).

²⁸ Boonin (2011: 336-7), Eidelson (2015: 177-87).

²⁹ Pace Lever (2005: 101-2).

³⁰ For further discussion of this case and its relevance to the racial profiling controversy, see Risse & Zeckhauser (2004).

That is because harms of this sort are likely to depend on specificities of context and/or the particular invidious stereotypes associated with particular groups.

For example, Randall Kennedy (1999) has argued that those who defend racial profiling neglect “the large extent to which racial profiling constantly adds to the sense of resentment felt by blacks of every social stratum toward the law enforcement establishment.” (33) Emphasizing how such resentment and alienation hampers the ability of the police and courts to adequately carry out their functions, he insists that “we need to be mindful of the deep reservoir of anger toward the police that now exists within many racial minority neighborhoods. Racial profiling is a big part of what keeps this pool of accumulated rage filled to the brim.” (33)

The importance of keeping in mind the significance of racial profiling in its relation to the historical and contemporary context of the profiled group is also evident in points raised by Annabelle Lever (2005). Lever argues that “[r]acial profiling publicly links black people with a tendency to crime” (97) and is therefore “likely to exacerbate the harms of racism” (97); it “is likely to remind blacks, all too painfully, that odious claims about their innate immorality and criminality justified their subordination in the past” (97) and “perpetuate, as well as reflect, white tendencies to draw invidious and complacent racial distinctions” (97).

Similarly, consider Deborah Hellman’s (2014) argument that racial profiling is wrong because it conveys a demeaning message directed at black Americans. In light of the cultural meaning of race in America, Hellman suggests that a policy of racial profiling expresses more than just the statistical generalization on which the policy relies: “Because its content tracks familiar cultural stereotypes about blacks and crime, the practice also expresses something more troubling such as ‘blacks are naturally disposed to be criminals.’” (237) In her view, “the cultural baggage of the category ‘black’ brings this meaning forward as well.” (237) Thus, Hellman similarly identifies the objectionable features of racial

profiling as contingent on specificities of history and context that apply selectively to certain racial categories in contemporary American society.

Harm-based arguments of this kind therefore seem able to explain straightforwardly why we are especially concerned about racial profiling that targets specific minority communities, such as black Americans. They also seem able to provide a straightforward explanation for why we aren't similarly concerned about incorporating race in suspect descriptions. A suspect description identifying the perpetrator of a given crime as belonging to some racial category does not of itself suggest any statistical regularity connecting members of that category to that type of crime, nor is it likely to elicit causal hypotheses about the basis of any such regularity. Incorporating race into a suspect description is therefore relatively unlikely to express, elicit, or reinforce demeaning stereotypes that cause stigmatic harm. And this may be thought to explain the important moral difference that we perceive between relying on race as a basis for criminal profiling and as an identifying characteristic in suspect descriptions.³¹

Harm-based objections to racial profiling therefore have considerable plausibility. But they also face a significant problem: arguments like those we've just considered do not seem well-placed to adequately condemn racial profiling that is practiced without public knowledge.³² Kennedy emphasizes the resentment and alienation engendered by racial profiling, which would seem to require that people know that profiling is going on.³³ Lever stresses the harmful effects that arise from publicly linking black people with a tendency to

³¹ Hellman (2014: 238-9), Eidelson (2015: 209-10).

³² Boonin (2011: 337-41).

³³ We might think that anger and resentment of the kind highlighted by Kennedy will be elicited simply in virtue of subjecting members of the black community to a higher absolute and/or relative level of policing, even if the fact that this is due specifically to racial profiling is unknown. I address this concern in the next two paragraphs.

crime. Presumably, no such public link is established if racial profiling is practiced covertly. Similarly, Hellman's emphasis on the insulting meaning expressed by racial profiling seems out of place if profiling is practiced in secret: it is unclear in what sense a demeaning message could be conveyed to black Americans by racial profiling that is carried out covertly.³⁴

It may be countered that even covert profiling will result in significant harm. Secret or not, racial profiling means increased stops targeted at black Americans. Such intrusions are likely to be experienced as threatening, demeaning, or humiliating even if the targeted individual isn't in a position to know whether they have been subject to racial profiling. In the worst cases, such stops result in lethal violence against black drivers, a fact chillingly reinforced by recent high-profile police killings whose victims have included Walter Scott, Samuel DuBose, and Philando Castile. We should be under no illusions, therefore, as to the seriousness of the harms that are likely to arise from subjecting black Americans to a higher level of scrutiny and police intrusion.³⁵

It remains unclear, nonetheless, that these harms will fully explain the strength and/or character of our opposition to racial profiling. That is because there are other policies that can lead to a higher level of stops, searches, or pat-downs suffered by black Americans, but that do not strike us as wrong to the same extent or at least not in the same way. For example, in a law enforcement environment in which police routinely rely on pretext traffic stops to fish for evidence of more serious crimes, a new law requiring drivers to signal when they switch lanes may expectedly increase the rate of traffic code violations by motorists and thereby also the level of pretext stops experienced by drivers, including black drivers. But except insofar as enforcement of the law is itself tainted by racial

³⁴ Cf. Lippert-Rasmussen (2014: 131), Segall (2012: 87).

³⁵ Cf. Lever (2005: 103-4).

profiling, we are unlikely to be strongly opposed to the adoption of a law of this kind in the same way that so many of us are opposed to racial profiling.³⁶

Eidelson (2015) suggests that the right response for proponents of harm-based arguments is, in effect, to concede the implication as regards covert profiling, at least in the following sense. Believing that the impermissibility of at least some cases of unalloyed racial profiling depends ultimately on the capacity of profiling to cause harms that require public knowledge, he concedes that unalloyed racial profiling could in principle be carried out permissibly if carried out in secret and might in some hypothetical circumstances where it could be carried out entirely harmlessly incur “no substantial moral objection ... at all” (217). He understandably describes this conclusion as “striking” (217).³⁷ As David Boonin (2011) notes, “virtually everyone who opposes racial profiling believes that profiling would be immoral even if it could successfully be kept secret” (339).

³⁶ Of course, nothing in the example as I have so far described it suggests that black drivers are disproportionately likely to be the target of this increased level of traffic stops. But we can add some further details to make this plausible. Suppose, for example, that a similar law has already been enacted in a neighbouring state, whereas black residents are less likely to have visited there because they are less likely to have relatives living there. As such, they have less prior experience in complying with the turn-signal law and are more likely to break it when it comes into effect. Even with these points in mind, I expect that we are unlikely to regard the higher absolute and relative level of stops suffered by black drivers in this case as an injustice of the same magnitude or kind as racial profiling. For further discussion of whether profiling is unfair by virtue of imposing higher relative burdens on black Americans, see Risse and Zeckhauser (2004: 156-160, 163-7), Boonin (2011: 341-7).

³⁷ To be clear, Eidelson (2015) “do[es] not believe there is a persuasive real-world argument to be made for secret racial profiling, even if one takes the decisive arguments against racial profiling to be those that rest on contingent consequences of the attitudes it engenders.” (216)

There is far more that could be said and has been said about the ethics of racial profiling in the recent philosophical literature,³⁸ but the foregoing discussion sets out what I take to be the principal dilemma arising from current debate on this issue. Harm-based arguments have significant plausibility in explaining why even unalloyed profiling is wrong when targeted at black Americans. However, harm-based arguments seem to encounter difficulties when we consider issues of publicity and secrecy. The view that racial profiling is intrinsically unjust faces no similar obstacle with respect to condemning racial profiling carried out in secret. But anyone who favours this view needs to explain what makes racial profiling directed at black Americans or other similarly situated groups as objectionable as it is generally regarded as being if not its characteristic harmful effects, given that we aren't strongly opposed to other forms of criminal profiling or statistical discrimination and, furthermore, accept that being black can in some cases be a contributing factor that leads to increased suspicion when police are searching for an individual matching a credible suspect description.

What I aim to do in the remainder of this paper is try to tackle the second horn of this dilemma.

4. Racial profiling as intrinsically unjust

This section aims to articulate a plausible moral principle, allowing us to conclude that racial profiling is *pro tanto* wrong in and of itself because it flouts a constraint on fair treatment. Articulating this principle doesn't allow us to make all that much progress in addressing the problems set out in the previous section, for reasons that I'll explain. It does

³⁸ For example, I have not discussed how racial profiling and police violence impact on the acceptability of framing racial inequality in terms of 'white privilege', as explored by Zack (2015), nor whether racial profiling of black Americans fails Cohen's (1992) interpersonal test of 'comprehensive justification', as suggested by Lippert-Rasmussen (2006) (see Risse 2007; Bou-Habib 2011; Eidelson 2015: 188-9).

allow us to see how we can view racial profiling as intrinsically problematic without becoming committed to a similar view about the use of race in casting decisions and suspect descriptions. But otherwise the present section is more a matter of laying the groundwork for the view about pattern-based reasons to be discussed in section 5.

The principle proposed here is intended to encapsulate and articulate the core moral intuition underlying the thought that individuals should be judged on the basis of their own conduct, rather than on the basis of generalizations about their race or gender. I will first present the principle and then explain how we are to understand it. Here it is:

- P1: It is *pro tanto* wrong for x to treat y worse than x treats or would treat z , if
- (a) x 's motivating reason for so treating y involves a merit-relevant judgment about y , where this judgment is based on x 's belief that y is F and
 - (b) (x has sufficient reason to believe that) y is not responsible for being F , and
 - (c) (x has sufficient reason to believe that) y is not responsible for the fact that her being F is used by x as a basis for so treating y .

This principle requires significant unpacking. For example, clause (a) restricts our attention to cases of differential treatment that rely on merit-relevant judgments. What do I mean by 'merit-relevant judgments'?

As I intend to use the term, a judgment that y is F is merit-relevant if and only if being F is contextually relevant for what treatment y merits or deserves. Thus, in a hiring context governed by meritocratic norms, judgments about skill-levels or talents represent merit-relevant judgments. In the context of typical legal proceedings, judgments of liability and guilt are merit-relevant. Similarly, in the context of police scrutiny and intrusion, judgments about whether someone is currently or prospectively engaged in criminal activity are merit-relevant.

I assume that when we speak of a person ‘meriting’ or ‘deserving’ something, we are not merely saying that their having it would be morally desirable. Following David Miller (1999: 133-8), I assume that judgments involving desert must invoke a particular kind of ground as a moral basis for distributing benefits and burdens: an act, omission, or sequence of acts and/or omissions on behalf of the agent, for which the individual is appropriately viewed as responsible. Responsibility, so understood, is not limited to moral responsibility: for example, it applies also in respect of the kind of praise that we can merit through the deliberate exercise of skills and talents, such as musical ability. Judgments involving desert invoke what Stephen Darwall (1977) calls *appraisal respect*, an evaluative attitude that we manifest toward others insofar as we judge them “either as a person or as engaged in some particular pursuit.” (38)

In the primary sense, then, a person is said to deserve something based on a performance that has already taken place, as when a student deserves a prize for having written the best undergraduate dissertation. In a secondary sense, the performance may be prospective: in this sense, we might say of a talented and hard-working team that they deserve to win, where this means that we expect them in the normal course of events to give a performance that would merit victory (in the primary sense).³⁹ Characteristics, such as innate physical beauty, which are borne by the agent in a merely passive sense, cannot be appropriate grounds of desert: they may attract wonder and amazement, but can’t entitle a person to appraisal respect from others.^{40 41}

³⁹ Miller (1999: 136-7).

⁴⁰ See Darwall (1977: 41-2), Miller (1999: 137-8).

⁴¹ We should keep in mind, of course, that our ordinary conception of individual responsibility extends beyond what we do (or forbear to do) deliberately, even if the latter figures centrally in our thinking about merit and desert. We can be blamed for indeliberate omissions, and some philosophers believe we can be blamed for such failures even when they cannot be indirectly traced to our deliberate activity, provided that they can be

Here is a final point of clarification about merit-relevant judgments. In focusing on cases where merit-relevant judgments are used to inform our treatment of others, I don't intend that we should commit ourselves to any general meritocratic theory of distributive justice. It's quite compatible with the view proposed here that merit-*ir*relevant judgments may be used in allocating jobs or other positions of advantage without anyone thereby being made subject to any unfairness. I'm convinced they may. The view being proposed is about a kind of unfairness that arises when merit-relevant judgments are relied on in guiding our treatment of others and drawing distinctions between people. It is not intended to tell us precisely when judgments of that kind ought to be used as a basis for distribution.

This point is important for the following reason. When a white actor is passed over in casting for the part of Ruth Younger because she doesn't match the race of the character, she may be thought to be disadvantaged in some respect on the basis of her race in a way that satisfies clauses (b) and (c) in P1. However, race is not relied on here in drawing merit-relevant judgments about the actor: her race is not used as a basis on which to appraise her talent as a performer. Therefore, clause (a) is unsatisfied. This sort of treatment need not involve any unfairness, therefore.

Some readers might suspect now that clause (a) is just an *ad hoc* modification designed to ensure that we reach this verdict about casting decisions. That isn't so. A clause

understood, roughly speaking, as expressive of what we value (or fail to value) deep down. See Levy & McKenna (2009: 115-118). Thus, in settling what it means for a trait 'to be borne in a merely passive sense' we may need to think in terms of whether this trait can be viewed as expressive of our fundamental cares, and not merely in terms of whether the trait and/or its exercise lies within the purview of our deliberate activity. More generally, I should emphasize that such views about responsibility as I invoke in this paper are meant to be philosophically non-committal so far as is possible and compatible with whatever turns out to be the best overall account of the nature of individual responsibility and appraisal respect. Thanks to an anonymous referee for taking me up on this point.

like (a) has to be included if we're to capture the thought that there is something problematic about judging people on the basis of their race or gender, as opposed to their own conduct. To *judge* a person, in this sense, is not merely to form a belief about her involving just any trait whatsoever. Except perhaps in certain unusual contexts, it would sound misplaced to allege that I have *judged* someone on the basis of her gender if I have inferred that she is likely to be the shorter of a pair of Americans about whom I know only that one is taller than the other and that one is female, the other male. Ordinarily, to make an inference about someone's height is not to judge her. To count as judging her, something stronger is needed. The proposed restriction to cases involving merit-relevant judgments is designed to answer to that need and to spell out just what it means to judge someone in the relevant sense.⁴²

Here is a different issue that needs clarification, which relates to clauses (b) and (c). Why did I insert parenthetical restrictions that invoke what the agent has sufficient reason to believe regarding *y*'s responsibility? Why aren't these clauses simply framed in terms of what *y* is in fact responsible for?

In formulating clauses (b) and (c) in this way, I'm motivated by the general presumption that our moral obligations are subjective in character, where this is understood not as a meta-ethical position, but as a first-order view according to which a person's obligations are a function of her beliefs, evidence, and/or credences about her circumstances, rather than the circumstances themselves.⁴³ Thus, I'd say that a person acts contrary to her moral obligations if she performs an action that she believes and/or ought to believe will result in many deaths, even if that action in fact saves many lives. A general defence of the

⁴² Cf. Wasserman (1991, 1997, 2011), who argues that reliance on naked statistical evidence is problematic to the extent that such evidence is used to draw inferences about morally significant behaviour.

⁴³ For views of this kind, see Jackson (1991), Prichard (1932), Ross (1939), Zimmerman (2006). For doubts, see Bykvist (2011), Graham (2010), Srinivasan (2015).

view that moral obligation is subjective in this sense lies beyond the scope of this paper, but I hope the view has sufficient independent plausibility that it can be allowed as an unargued assumption.

With these points in mind, I believe we're now in a position to see P1 as articulating the following kind of moral ideal. In our treatment of others, we should respect the significant interest that we each have in controlling how we are judged by others, such that we can be deemed appropriately responsible for the respect and/or disrespect that we are shown by the people around us by virtue of how we are treated. Hence, we act in a way that's *pro tanto* wrong if, in drawing distinctions between individuals, we judge them on the basis of certain traits and we have sufficient reason to believe that they can neither be held responsible for those traits nor for the fact that they are so judged in light of them. Doing so shows insufficient respect for the interest that we each have in being appropriately appraised in light of our own conduct.⁴⁴

That's all I'm going to say in terms of clarifying P1. Let's now consider its implications for the issue at hand.

P1 straightforwardly entails that racial profiling is *pro tanto* wrong. Intuitively, judgments about criminal behaviour are merit-relevant with respect to police scrutiny and intrusion: in the context of deciding who to arrest or detain, judgments about whether

⁴⁴ Should we regard a person as capable of being wronged merely in light of how she is judged, even if these judgments remain private and do not inform how she is treated? It certainly seems intuitively plausible that judging another person on the basis of her race or gender can be of itself discreditable; I think we are apt to condemn those who view others through the lens of racial stereotypes even when nothing otherwise comes of this. But some may believe that we cannot specifically *wrong* another person merely in light of our private attitudes, even if such attitudes may cast *us* in an unfavourable light. The case for thinking that a person is wronged is much clearer if we assume that such judgments inform the manner in which she is treated, and I therefore focus on these instances in my discussion. For more on the question of how attitudes and behaviours interact in our assessment of wrongful discrimination, see Arneson (2006: 780-3), Eidelson (2015: 76-84).

someone is currently or prospectively engaged in criminal activity are relevant in respect of whether a person merits such treatment. We are all in a position to know that no one is responsible for her race and, similarly, that no one can realistically be regarded as individually responsible for the fact that her race is a statistical predictor of crime, if it is. Each of the conditions for unfair treatment specified in P1 is satisfied. This holds regardless of whether racial profiling is or is not a matter of public knowledge. My view thus identifies a particular wrong-making feature associated with racial profiling *per se*.

By contrast, P1 does not entail that it is *pro tanto* wrong to use race or ethnicity as a factor when searching for an individual matching a suspect description. Here is why. Suppose that someone has been detained as a suspect in a mugging in part because he is South Asian and therefore matches the description of the attacker offered by the victim. The use of ethnicity as a factor in detaining the suspect entails that clauses (a) and (b) are satisfied. But provided that the police do not already have evidence that rules him out as the guilty party, condition (c) is not satisfied. Why not? Well, if the suspect *is* the mugger, then it seems reasonable to suppose that he *can* be held responsible for the fact that police are seeking a South Asian individual. Therefore, provided he cannot already be ruled out as a suspect on the basis of further evidence, condition (c) is unsatisfied. The principle I've proposed thus preserves our intuitive sense that there is an important asymmetry between racial profiling and the use of race in suspect descriptions. The key difference between the two turns on the manner in which the discriminator may be deemed to knowingly disregard the responsibility of the discriminatee in drawing inferences about her involvement in criminal activity when engaged in racial profiling, but not when matching a person to a suspect description.

As advertised, P1 allows us to view racial profiling as intrinsically problematic in a way that doesn't conflict with our intuitions about the use of race in casting decisions and suspect descriptions. In what respect, then, does it fail to make significant progress on the

problems noted in section 3? It fails to do so because P1, considered of itself, does not allow us to satisfactorily explain why racial profiling of black Americans differs in a morally significant way from the other cases of statistical discrimination noted in the previous section, where unchosen traits serve as a basis for disparate treatment and yet we find the practice acceptable.

The principle is perhaps not a complete failure in this respect. It could be that P1 can accommodate intuitions about some of these cases because they do not satisfy clause (a): *i.e.*, because they do not involve merit-relevant judgments. Consider, for example, affirmative action programs like those discussed in section 3. Suppose that an applicant's race is considered as a reason to admit her in light of the expectation that doing so will help to break down stereotypes and promote cross-racial understanding, both legitimate educational goals of the university. Does the university thereby rely on race as a basis for merit-relevant judgments about its candidates? Clearly it sees race as a relevant factor and as constituting some form of qualification, but need it do so in a way that differs altogether from the casting director searching for her Ruth Younger? Perhaps not. If not, (a) is unsatisfied.⁴⁵

Aside from complications of this sort, it seems clear that P1 is not going to help us in distinguishing racial profiling from acceptable acts of statistical discrimination like those discussed in section 3. Except insofar as we are able to quibble about satisfaction of clause (a), P1 is going to end up condemning all of these cases as involving the same kind of wrongful treatment. But, in my view, this is not a defect in the principle. It only tells us that

⁴⁵ Notably, for critics such as Pojman (1998) a key problem with affirmative action policies is supposed to be that they go against a meritocratic allocation of university places. It may therefore be surprising that on the present analysis a failure to rely on merit-relevant judgments in implementing affirmative action policies could be viewed as morally desirable. But see Wasserman (1997) for a roughly similar view on the matter.

the principle is incomplete as a guide to understanding the significant wrong that occurs when racial profiling is targeted at groups such as black Americans.

Why do I say this? Keep in mind that P1 only tells us whether some action is *pro tanto* wrong, not whether the action is wrong all things considered. Therefore, P1 doesn't require us to say of any of these cases that they involve morally prohibited behaviour. In this respect, its implications don't contradict our intuitions, at least not insofar as the latter tell us simply that the relevant practices are permissible. The implications of P1 contradict our intuitions only if the latter tell us that there isn't *any* degree of unfairness involved in these cases. A view of that kind does not strike me as especially compelling, however.

Consider the case of gender discrimination in car insurance pricing. Suppose Frank is charged a higher price for insurance than Maya, not because the insurer has any specific evidence relating to his driving abilities or risk preferences, but simply because the insurer regards Frank as more likely to drive dangerously in light of his gender. There seems to me to be *some* element of unfairness involved here.⁴⁶ Obviously, other considerations can serve to explain why the practice is permissible in an all-things-considered sense. Still, it seems to me there needn't be any absurdity in supposing that male drivers are treated unfairly to some extent by this practice.

I'd want to take a similar line with respect to the other cases of statistical discrimination discussed in section 3. Assuming that these cases satisfy (a)-(c), they are cases where people are made subject to some degree of unfairness but where this is justifiable in an all-things-considered sense. I believe this can be said without serious cost to

⁴⁶ Gender discrimination in the pricing of car insurance has in fact been illegal in the EU since the end of 2012 (Collinson 2012). It is prohibited outright in nine US states, and strongly limited by law in many others (Avraham et al. 2014).

the overall plausibility of P1, either because it's intuitively plausible that some element of unfairness is in play or at least unobvious that it isn't.⁴⁷

My position, then, is this. The fact that P1 seems to count the various cases of statistical discrimination that we're considering as involving the same kind of unfairness doesn't indicate that we need to revise the principle. It is not absurd to believe that a similar kind of unfairness occurs in each of these cases. What is implausible is to suppose that any of the other cases we've discussed involve a wrong of the same degree of seriousness as that involved in racial profiling targeted at black Americans. The latter seems far graver and far more troubling. What P1 cannot explain is this difference in magnitude. If we are not going to rest this explanation on the harms caused by instances of racial profiling, is there any other way in which we can explain this difference in moral seriousness?

5. Racial profiling and collective wrongdoing

In order to explain why there is something especially objectionable about racial profiling targeted at black Americans and other similarly situated groups, we can appeal to the idea that an action that may be relatively innocuous considered in isolation may be rendered significantly morally problematic (and even wrong *simpliciter*) insofar as it forms part of (or may be expected to form part of) a pattern of similar actions that collectively bring about or realize a very bad or morally catastrophic outcome.

This way of thinking gains support from consideration of cases such as the following. Suppose that if everyone drives their own car to work, the air quality in the city will deteriorate considerably. If we all reliably cycle, walk, or take public transport, the air quality will remain very high. A significant decline in air quality depends on there being

⁴⁷ For a recent defense of the view that affirmative action involves a justifiable element of unfairness see Blum (2013). See also Nagel (1973), Ezorsky (1991: 81-4).

very many people who each drive to work regularly. Holding fixed the actions of every other commuter, the decision of a single individual to switch from cycling to driving makes no noticeable difference to the experienced air quality. In this case, it seems plausible that we each have non-trivial moral reasons to avoid driving our own cars to work. The weight of these reasons seems to be a function not of the negligible harm that we would each individually contribute via our emissions, but rather the harm that we together would do if sufficiently many of us decide to drive to work every day: we have reason to avoid performing actions that have the potential to form part of that sort of pattern.

Similarly, consider a variant of the ‘harmless torturers’ case discussed by Derek Parfit (1984: 80–82). Suppose a single torture victim is tied to a machine that provides electric shocks of a certain voltage. The machine is initially set at such a low voltage that the victim feels no pain. Turning a dial on the machine increases the voltage very slightly: so slightly that each adjacent setting feels the same to the victim. If the dial is turned sufficiently many times, the change in voltage will become noticeable and eventually the shocks will become unbearable. Imagine a very large group of torturers, each of whom turns the dial just once, resulting eventually in terrible suffering for the victim. We’re inclined to say that the torturers have acted wrongly. But since each through her action seems to make no difference to the level of discomfort suffered by the victim, it seems we must appeal to the cumulative effect of these actions in order to explain why each torturer has acted impermissibly.

Cases like these may be taken to support a general principle of the following kind:⁴⁸

P2: There is greater reason for x not to ϕ in C to the extent that (x has reason to believe that)
 the act of ϕ -ing in C will form part of a pattern of similar actions— y ϕ^* -ing in C^* , z ϕ^{**} -ing

⁴⁸ For this sort of view, see Andreou (2014), Kutz (2000: 166–203), Parfit (1984: 70–73, 75–83), Spiekermann (2014), Woodard (2008).

in C^{**} , *etc.*-such that these actions together bring about or realize a very bad or catastrophic outcome, O .

It's by no means uncontroversial that the cases we've just considered support P2. Some philosophers doubt that these cases are well-described: they are sceptical of the possibility that each one of many actions considered individually could make no difference whereas the many actions considered collectively can.⁴⁹ I believe this is possible and that we need to endorse P2 to accommodate our intuition about cases like that of the harmless torturers.⁵⁰ But as I don't have the space to defend this view here, philosophers who reject P2 will have to treat my argument in what follows as establishing only a conditional conclusion in which P2 figures in the antecedent.

The basic outline of the proposal to be considered here is then as follows. If a law enforcement officer subjects a white individual to racial profiling in the course of attempts to apprehend a serial killer like the Beltway sniper and a black individual is subjected to racial profiling by a law enforcement officer attempting to apprehend drug couriers, both individuals are treated unfairly by virtue of being made subject to the same kind of wrong, as identified by P1. Considered in and of itself, this form of unfairness is relatively minor, as befits our intuition that it is relatively unproblematic to profile for white suspects in attempting to apprehend a serial killer. Even if these acts are otherwise the same in other morally relevant respects, the wrong involved in detaining the black individual is much

⁴⁹ On this issue see Arntzenius and McCarthy (1997), Fara (2001), Kagan (2011), Keefe (2011), Nefsky (2011), Parfit (1984: 78-82), Voorhoeve and Binmore (2006).

⁵⁰ Even if it is not possible, there are other cases that seem to support P2 but involve overdetermination, pre-emption, or coordination problems, rather than cumulative harm. See Parfit (1984: 70-3), Woodard (2008: 42-7).

greater by virtue of the fact that it forms part of a larger pattern of racial injustice that permeates American society.⁵¹

How exactly should we understand the larger pattern involved here? One possibility would be to see particular acts of racial profiling as part of a sequence of injustices that collectively restrict and burden the profiled individual. As a victim of racial profiling, Dr. Elmo Randolph was stopped more than 50 times by New Jersey state troopers commuting to and from work between 1991 and 2000, despite never being ticketed.⁵² Even an isolated stop may be construed as part of a broader pattern of racially biased unfairness to which a person is made subject over the course of his or her life.

This approach faces the following challenge, however. Consider a black man who is made subject to racial profiling and also faces additional statistical discrimination, being charged more for auto insurance on account of being male. It seems somewhat odd to say that the insurer has greater reason to avoid statistical discrimination against this individual on account of his gender, as opposed to an otherwise similar white customer. At the same time, we want to be able to say that law-enforcement officers are party to a more severe form of injustice if they racially profile blacks as opposed to whites. Both acts involve statistical discrimination and may be thought to involve a similar kind of unfairness visited upon the same individual in the case we're now considering. Suppose we grant that the unfairness in racial profiling is intensified by virtue of being part of a larger pattern. Why is the unfairness of statistical discrimination on the basis of gender in the insurance market not intensified in the same way? Why is it not part of the same pattern?

In one respect, of course, the answer is obvious. The unfairness in racial profiling is conditional on race and can therefore be grouped with a larger sequence of race-based

⁵¹ Cf. Lippert-Rasmussen (2014: 169-70) on discrimination and cumulative harm.

⁵² Ramirez et al. (2000: 5).

injustices. Although it is likely that the person has experienced discrimination on account of prejudices about black men,⁵³ there is plausibly no comparable sequence of injustices suffered, as in this case, in virtue of being a man *tout court*. But this only raises the question of why we should group harms in this way. If we see particular acts of racial profiling as part of sequence of unfair disadvantages that collectively burden the profiled individual, why is the unfairness involved in being charged more for car insurance on account of one's gender not part of the sequence? In searching for the cumulative injustice of which racial profiling forms part, it's not altogether clear why it should matter in and of itself which dimension of identity is the ground of unfair treatment, assuming our focus is simply on the individual and on individual disadvantage.

It may well be that others are able to see a solution to this problem that I have overlooked. Nonetheless, the view that I'll explore here asks us to see racial profiling of black citizens as especially problematic insofar as it is one of a number of injustices that together constitute racial oppression, where oppression represents an injustice that pertains, in the first instance, to the overall standing of a group of persons.

Within the conceptual framework developed by philosophers including Marilyn Frye (1983), Iris Marion Young (1990), and Ann E. Cudd (2006), 'oppression' denotes a systematic pattern of injustice suffered by members of a socially salient group in virtue of their membership within that group and through which that group is subordinated to others within society. Oppression, so understood, involves a pattern of injustice embedded within diverse dimensions of social life. This stipulation is typically paired with the observation that this overarching pattern need not represent the deliberate aim of any person. Oppression is conceived as something that arises to a large extent through the texture and

⁵³ Stereotypes about black men are obviously an important factor in explaining police violence: see Zack (2015: 22-3). On the interaction of race and gender in labour market discrimination, see Lang & Lehmann (2012).

patterns of ordinary life, the cumulative product of very many individual actions that may appear relatively minor if viewed in isolation.⁵⁴ In this vein, Frye (1983) memorably compares oppression to a bird cage: 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 other key element in the contemporary understanding of oppression is that it is taken to be a condition pertaining in the first instance to the standing of social groups, with individuals qualifying as oppressed in virtue of their membership within these groups.⁵⁵ In the primary sense, then, oppression is a systematic pattern of injustice suffered by members of oppressed groups *qua* members of those groups: for example, women *qua* women or blacks *qua* blacks. As such, not every injustice suffered by a member of an oppressed group constitutes an element of the relevant oppressive structure, but only those explained by his or her membership in the relevant group: *i.e.*, a group the members of which are subject to systematic injustice in virtue of their membership within said group. Oppression, so understood, is a form of injustice that involves the subordination of one group to another and contributes to caste-like divisions within society.

That having been said, we need not think of oppression as a harm that pertains to the social group considered as some kind of supra-individual moral patient. A view that asks us to see oppression as a paradigmatic form of injustice is quite compatible with the assumption that the standing of the group matters only insofar as it matters to its

⁵⁴ See Frye (1983: 376-8), Kernohan (1998: xi, 11-13) Young (1990: 40-42). Obviously, this is not meant to exclude that violence and other similarly serious wrongs can figure as significant elements in oppressive social structures: see Young (1990: 61-3).

⁵⁵ Cudd (2006: 23-54), Frye (1983: 379-80), Young (1990: 42-8).

members.⁵⁶ Membership in many socially salient groups represents a source of identity, community, and self-understanding. Under such conditions, individuals characteristically feel insult or indignity at injustices that threaten their standing. The fortunes of the group as a whole may come to represent one important component of individual self-respect. This important interest that individual people take in the status of identity-forming groups is clearly frustrated wholesale when such groups are made subject to oppressive societal conditions. And, quite obviously, oppression is associated with a range of other serious harms, being composed of injustices suffered by individuals in virtue of their group membership and tending to perpetuate and reproduce such harms by disempowering disadvantaged groups within the political and cultural spheres.⁵⁷

The key relationship between oppression and the especial wrong involved in racial profiling targeted at black Americans is then as follows. Like air pollution, oppression represents the cumulative effect of many individual actions whose moral seriousness we underestimate unless we adopt a macro perspective that places such actions within a larger pattern of similar actions.⁵⁸ It is by placing the racial profiling of black Americans within such a pattern and appealing to P2 that we understand why the unfairness inherent in racial profiling of any kind is made significantly worse when acts of racial profiling are targeted at members of this group or other similarly situated groups, quite apart from their capacity to generate such characteristic harmful effects as stigmatization or a breakdown in trust between police and profiled communities.⁵⁹

⁵⁶ Cudd (2006: 52-3).

⁵⁷ See Khaitan (2015: 117-39), Young (1990: 48-63).

⁵⁸ See Kernohan (1998: 9-13) for the comparison of oppression and pollution.

⁵⁹ We can arguably imagine instances of covert unalloyed racial profiling that are entirely harmless: see Eidelson (2015: 216-7). Is it really plausible that such cases also form part of this pattern? Might only harmful injustices count as such? I am grateful to an anonymous referee for pressing me on this issue. I think it's clear

This approach allows us to explain quite straightforwardly why there is no similar intensification of the wrong-making unfairness inherent in statistical discrimination when a black man is charged a higher price for car insurance on account of his gender, something that seemed difficult to explain on a purely individualist construal of the cumulative harm of which racial profiling forms part. As noted, not every injustice suffered by the members of an oppressed group represents an element of their oppression: it's only those that befall such individuals by virtue of their membership within the relevant group that count as such. Unfairness that someone suffers only in virtue of being male and not even partly in virtue of his race therefore cannot be understood as a constituent in racial oppression.⁶⁰

that even harmless injustices should be counted. Consider an apparently harmless instance of prejudiced direct discrimination. Suppose someone is denied a job at one firm due to racist bias and so takes up a job at a different firm instead. This may not sound harmless, but additional circumstantial details may seem to make it so. The other firm is, in fact, a much better place to work. There are no other costs associated with taking up a job at the other firm: for example, this doesn't require a longer waiting period. We can even suppose that when the applicant learns that their application was disregarded on racial grounds, they merely laugh this off as absurd. In this case, it may be difficult to say that the applicant is harmed. At least, it may be difficult to do so without appeal to a specialized conception of discriminatory harm that will likely also put a serious question mark over the possibility of genuinely harmless racial profiling. (See Lippert-Rasmussen 2014: 157-162 on the options available for defining discriminatory acts similar to that discussed in this footnote as harmful.) Even so, if someone had been gathering evidence that showed members of this group to be subject to oppression, I think we would find it very odd if they were to disregard this instance of discrimination and treat it as falling outside the scope of the phenomenon that interests them.

⁶⁰ When cumulative injustice is suffered in the form of intersectional discrimination, should such injustices be thought of as forming their own individual pattern, or should they be grouped as part of the broader patterns that pertain to the individual constituent characteristics that define such groups? Or should it be both? See Khaitan (2015: 162-3) on this question. This issue requires greater attention and exploration in working out the details of the approach proposed in this paper.

The view discussed here exhibits obvious affinities with the position advanced by Owen Fiss (1976), Catharine MacKinnon (1979), Cass Sunstein (1994), and Tarunabh Khaitan (2015), according to which the primary justifying goal of anti-discrimination law should be thought of as the elimination of systematic and significant group-based disadvantages that result in a caste-like society. There is an especially strong resemblance to Khaitan's view that discriminatory conduct in its paradigm instances is "doubly wrongful" (197) because it not only "makes a person suffer some adverse effect owing merely to her membership of a normatively irrelevant or even valuable group" (197)⁶¹ but also "exacerbates the existing advantage gap between groups" (197). In a similar fashion, the position set out here is that racial profiling targeted at black Americans or other similarly situated groups involves an especially serious injustice because it violates P1 *and* falls within the scope of P2. But the view proposed here avoids some objections that seem to apply to Khaitan's theory.

In particular, Khaitan looks to have trouble when it comes to the use of race in casting decisions and suspect descriptions. This seems to involve making a person suffer some adverse effect due to membership in a "normatively irrelevant" group⁶² and would seem thereby to exacerbate the advantage gap between groups insofar as the burdened individual belongs to a protected group. At the very least, there's a challenge here for Khaitan to explain why his view doesn't strongly condemn the use of race in casting decisions and suspect descriptions.

I believe the view set out here faces no similar challenge. It doesn't count these practices as involving any inherent injustice, since they don't violate P1. Nor would they fall within the scope of P2 by virtue of forming part of a pattern of similar actions that together

⁶¹ A normatively irrelevant group is here understood to be a group membership in which "should not affect how successful our lives are" (Khaitan 2015: 56).

⁶² For the definition of a normatively irrelevant group, see the previous footnote.

bring about racial oppression. Because ‘oppression’ denotes a pattern of pervasive injustice suffered by the members of certain groups in virtue of their group-membership, the inclusion of some action within this pattern requires that it be classifiable as a racially biased injustice prior to the point at which a macro-level, pattern-based perspective is invoked. Assuming P1, racial profiling is classifiable as such, whereas we are not required to view the use of race in casting actors for *A Raisin in the Sun* or in searching for a suspect matching a witness description that includes racial identity in this way.⁶³

6. Objections

In this section, I’ll consider some objections that might be raised against the views set out in this paper. Unless otherwise indicated, these are posed in the voice of an imagined interlocutor.

The overarching position set out in this paper seems self-contradictory. In section 4, you emphasize that we ought to be judged on the basis of our own conduct. However, in section 5, it seems you want to judge people in terms of the conduct of others, since you suggest that an action can be wrong not because of what is done by the agent herself but by the group to which she belongs. It seems as if you want us to be judged as individuals and as members of groups. Isn’t there a contradiction here?

⁶³ To say this of the use of race in casting a typical production of *A Raisin in the Sun* is not to imply, of course, that casting decisions pertaining to race cannot be morally criticisable just so long as they satisfy P1. For example, casting actors in a major blockbuster in a way that foreseeably reinforces damaging racial stereotypes represents a significant wrong, even if the manner in which the individual actors are chosen does not treat their race as a basis for making merit-relevant judgments about their abilities. Similarly, there are ways in which to handle the use of race in suspect descriptions that may be wrong, even if reliance on race as an element in a suspect description is not wrong *per se*. I’m grateful to an anonymous referee for pressing me to clarify this point.

I believe any contradiction here is merely apparent. Note, in particular, that P1 does *not* entail that a person can't be judged in light of her membership within a certain group or category. It leaves open that we can be judged in light of our membership in some group insofar as we can be viewed as personally responsible for belonging to the group and/or for being judged in light of our group membership. Furthermore, whereas there is a sense in which P2 asks us to judge people in light of what others do and not merely in terms of what we ourselves bring about, it does so only insofar as the relevant agent has reason to believe that her action will make her party to a group that collectively brings about a very bad or catastrophic outcome. As such, P2 plausibly falls within the aforementioned space left open by P1: it does not ask us to judge a person in light of her membership in some group except insofar as she can be viewed as responsible for being judged in light of her group membership.

P1 seems vulnerable to an objection raised by Matt Cavanagh (2002: 192-3) who considers a similar 'responsibility principle' as an objection to statistical discrimination by race in allocating advantageous positions. Cavanagh argues that it is strange to insist that people should be judged on the basis of conduct for which they can be held responsible given that assessors are very often ultimately interested in predicting traits for which candidates or applicants can't be held responsible, such as their level of intelligence. A principled stance requires applying the 'responsibility principle' to what assessors are ultimately interested in, Cavanagh argues. But are we willing to say that universities and employers shouldn't select people based on their intelligence or that doing so constitutes the same kind of wrong as statistical discrimination based on race?

I believe this objection can be seen to rest on a misunderstanding of what employers and university admissions officers are interested in when they select candidates based on their

perceived intelligence. What they are ultimately looking for is not someone who passively exemplifies a given level of intelligence, but someone who actively exercises and will exercise her intelligence in solving certain problems. It needn't be implausible to suppose that a person can be deemed appropriately creditable for what they accomplish in deliberately exercising their intelligence in this way. Granting this assumption, the mismatch suggested by Cavanagh is illusory, since what assessors are ultimately interested in doesn't lie beyond the purview of individual responsibility.⁶⁴

An anonymous referee writes: "The case of the harmless torturers looks relevantly different from racial oppression: in the latter but not the former, the contributors vary significantly in their level of contribution to the collective harm. Why should we count the negligible contributions as significant wrongs, rather than counting only contributions of or over a certain magnitude? Or allocate the wrongs proportionately to the harms risked by the individual's action? It is not obvious that the (contentious) conclusions from [the] harmless torturers case can be straightforwardly applied to oppression."

It is true that the harmless torturers case, as traditionally described, involves a range of individuals who each make the same contribution. However, I don't believe that this has any deep moral significance in terms of what we learn by reflecting on this case.

Suppose we re-imagine the case, as follows. There are very many who turn the dial only once, in such a way that the victim cannot feel any difference. But there are also some who turn the dial across a number of steps, such that there is a noticeable increase in the suffering of the victim. Suppose that some choose freely to turn the dial across many steps

⁶⁴ My reasoning in responding to this objection closely follows Miller (1999: 143-9).

in this way, but that the cumulative impact of the single-dial turns also contributes substantially to the overall pain that the victim eventually feels.

When the victim ends up in agony as a result of all the turns effected on the dial, I think we would still regard those who had turned the dial only once as having acted wrongly. Of course, we would assign greater blame to those who had turned the dial across many steps. In that respect, I agree entirely with the suggestion that wrongs should be allocated proportionately. Even so, I think we would continue to condemn those who had turned the dial only once. On this basis, I deny that the difference highlighted by the referee is morally significant. Cases involving variability of contributions can also support the view that our reasons against performing certain actions may be stronger where such actions are expected to form part of a pattern of similar acts that together bring about a very bad outcome, such that acts that might appear of minor moral significance when viewed in isolation take on the character of serious wrongs when viewed as part of such a pattern. Such cases also fall within the scope of P2 and may be relied upon in illustrating the intuitive plausibility of that sort of principle.

*Because criminal victimization tends to be predominantly intraracial, black Americans are disproportionately likely to be the victims of those crimes that are perpetrated by other black Americans. To the extent that unalloyed racial profiling of black Americans might represent an effective crime prevention strategy, one might then infer that it will disproportionately benefit members of the black community. How, then, could it be said to contribute to racial oppression?*⁶⁵

To add force to this worry, let me note the following. It hasn't been my aim in this paper to show that racial profiling cannot be justified in any possible world. If racial profiling can in

⁶⁵ Cf. Risse and Zeckhauser (2004: 162-6), Boonin (2011: 341-2).

fact substantially lower the rate at which people are made victims of serious crime, it's clear that there are strong moral reasons favouring the adoption of such a policy. The question is how to balance such reasons against those arising in light of the serious costs. This paper has been trying to add to our understanding of the nature and strength of the reasons that speak against racially profiling black Americans or members of similarly situated groups. Being generally opposed to moral absolutism, I am nonetheless open to the possibility that the reasons in favour of racial profiling could outweigh the reasons against in some possible world.

Imagine, then, that we are in the nearest possible world of that kind and let's grant the objector that the benefits of profiling are here concentrated within the profiled community. Could we still plausibly maintain, in that case, that there are especial reasons against racial profiling, owing to the fact that this form of racially biased unfairness constitutes an element of racial oppression?

I believe so. There is nothing incoherent in supposing that actions that contribute to the oppression of a social group may also serve to benefit that group, not even where 'benefit' is understood in an all-things-considered sense. Actions and policies have many effects, direct and indirect, some of which may counter-balance one another. For example, acts of serious injustice perpetrated against certain groups can serve to strengthen their resolve and political will, leading to social changes that significantly improve their lives overall. These subsequent improvements don't undo the fact that the triggering injustice formed part of an oppressive social environment. There is nothing implausible, therefore, in the thought that acts or policies that contribute to oppression can in some possible cases also serve to benefit the oppressed group.⁶⁶

⁶⁶ An anonymous referee raises the following question. If racial profiling were effective in reducing intraracial criminal victimization and served overall to benefit the profiled community, should we then view such a policy as akin to certain kinds of affirmative action policies in respect of being a policy that is *pro tanto* unfair (by

You say that P1 is intended to give expression to the significant interest that we each have in controlling how we are judged by others. However, the principle apparently fails to identify as problematic certain actions that seem to frustrate that interest. Consider someone who is detained by airport security because her behaviour matches the profile of a drug courier: she pays in cash, checks no hold luggage, etc., all for perfectly innocent reasons. Because her behaviour is voluntary, it seems that clause (b) in P1 is unsatisfied. But let's assume, plausibly, that she was unaware that her actions would render her a suspect in this way and had no way of knowing, since the profile is secret. Given her ignorance, it seems that her treatment fails to accord her meaningful control over how she is judged by others. To the extent that P1 is intended to honour the interest that we have in exercising such control, shouldn't it condemn this kind of criminal profiling as well?⁶⁷

I am not sure of the force of this objection. I agree that the passenger's ignorance of the profile significantly limits her ability to control how she is appraised in this context. However, it feels intuitively unclear that it does so to such an extent that her treatment is thereby rendered *pro tanto* wrong. After all, she is still judged, and treated accordingly, in light on how she herself chooses to act. It cannot reasonably be demanded of others that

virtue of relying on statistical discrimination grounded in racial criteria) but justified all-things-considered by virtue of its capacity to ameliorate group-based disadvantages? As the referee notes, an affirmative answer would suggest that the reasons against racial profiling that I highlight here would also apply to such affirmative action policies. But the salient point of contrast between the two policies is clearly that insofar as affirmative action involves statistical discrimination that falls foul of P1, it is the members of advantaged groups who are dispreferred on account of their racial background who are subject to some degree of unfairness on that account, and not those minority candidates who may be favoured by the policy. By contrast, the injustice inherent in racial profiling is one that also falls on the members of such disadvantaged minority communities.

⁶⁷ I am grateful to Ben Eidelson for raising this concern.

they should provide us with total control in how we are judged by them when such judgments inform how we are treated. It is difficult to say how much control we can reasonably demand and what deprivation of control we can count as sufficient to wrong us.

To the extent that we do find the objection convincing, the natural response would be to revise P1 so as to omit clause (b), leaving (a) and (c) in place. In the sort of case we are considering here, both clauses would be satisfied. The suggested revision also has the virtue of simplifying P1. And there may be additional reasons that favour this revision.

In particular, we may believe that it is *prima facie* desirable for an account of what is wrong with racial profiling to be straightforwardly extendible to cases of religious profiling. And it may be felt that P1, as it stands, provides an insufficiently clear basis from which to infer that profiling that relies on religious affiliation rather than race involves an inherent injustice, since a person's religious identity may be thought of as something for which she is responsible in a way that voids clause (b).⁶⁸ To me, the view that people are typically responsible for their religious identity seems far from clear, but the matter is complex and is likely to turn on difficult questions about the nature and basis of religious identity, as well as questions about the nature and purview of individual responsibility. If we do conclude that profiling based on religious identity cannot be captured by P1 as it stands and also believe that an account of the wrong in racial profiling should be able to generalize straightforwardly to cover such cases—and neither of these claims strike me as unreasonable—then we would have additional reasons to delete clause (b) from P1.

Of course, omitting clause (b) does also require us to recognize ordinary behavioural profiling as *pro tanto* wrong and as unfairly burdening those who are subject to heightened scrutiny, at least in those cases where the person has no reason to expect that her behaviour

⁶⁸ I am grateful to Ben Eidelson and Tom Sinclair for raising this issue, and to an anonymous referee at this journal for pressing me on it further.

will elicit increased scrutiny and intrusion. But this does not strike me as especially problematic. It does not sound altogether odd to suppose that behavioural profiling that we know will ensnare some innocent people on the basis of criteria they could not anticipate is unfair in some respect. We may think of this as yet another instance of small-scale unfairness that is justifiable all-things-considered. We should simply say, as we have said in other cases, that this sort of unfairness fails to merit especial moral concern because it does not form part of a broader and more pervasive pattern of similar injustices suffered by the targeted individuals in light of their membership within the profiled category. It may be that it was a mistake all along to think that a principle like P1 was crucial to distinguishing between racial profiling and other less controversial forms of criminal profiling, as opposed to letting principle P2 do the crucial work.

7. Summary and conclusion

By way of recap, the broad contours of the account presented here are as follows. Racial profiling violates a constraint of fairness that proscribes treating some individuals as meriting disadvantageous treatment where such treatment depends on facts about the disadvantaged subject that we have sufficient reason to believe lie beyond their responsibility. This kind of unfairness need not be all that serious considered on its own terms. However, when this kind of unfair treatment is visited on members of oppressed groups in light of their membership in such groups, this inherent injustice is exacerbated by virtue of figuring as a contributing element in a larger cumulative wrong. In this way, we can understand why it must be especially problematic for law enforcement officers to engage in racial profiling targeted at black Americans and other similarly situated groups. Furthermore, we can do so without having to appeal to harmful effects that depend on public knowledge, allowing us to recapture the sense of pre-theoretic opposition that so many of us feel when it comes to racial profiling.

References

- Andreou, Chrisoula (2014) The good, the bad, and the trivial. *Philosophical Studies* 169(2), 209–225.
- Arneson, Richard J. (2006) What is wrongful discrimination? *San Diego Law Review* 43(4), 775–808.
- Arntzenius, Frank and McCarthy, David (1997) Self torture and group beneficence. *Erkenntnis* 47(1), 129–144.
- Avraham, Ronen, Kyle D. Logue, and Daniel Schwarcz (2014) Understanding insurance antidiscrimination laws. *Southern California Law Review* 87(2), 195–274.
- Blum, Lawrence (2013) Racial and other asymmetries. In Deborah Hellman and Sophia Moreau, eds. *Philosophical foundations of anti-discrimination law*, 182–200. Oxford: Oxford University Press.
- Boonin, David (2011) *Should race matter? Unusual answers to the usual questions*. Cambridge: Cambridge University Press.
- Bou-Habib, Paul (2011) Racial profiling and background injustice. *The Journal of Ethics* 15(1–2), 33–46.
- Brennan, Samantha (2016) The moral status of micro-inequities: in favour of institutional solutions. In Michael Brownstein and Jennifer Saul, eds. *Implicit bias and philosophy, vol. 2: moral responsibility, structural injustice, and ethics*, 235–254. Oxford: Oxford University Press.
- Bykvist, Krister (2011) How to do wrong knowingly and get away with it. In Rysiek Sliwinski and Frans Svensson, eds. *Neither/nor: philosophical papers dedicated to Erik Carlson on the occasion of his fiftieth birthday*, 31–47. Uppsala: Uppsala University.
- Carlson, Darren K. (2004) Racial profiling seen as pervasive, unjust. *Gallup*, July 20th 2004. <<http://www.gallup.com/poll/12406/racial-profiling-seen-pervasive-unjust.aspx>> Last accessed July 18th 2017.
- Cavanagh, Matt (2002) *Against equality of opportunity*. Oxford: Oxford University Press.
- Cohen, G. A. (1992) Incentives, inequality, and community. *The Tanner Lectures on Human Values* 13, 263–329.
- Collinson, Patrick (2012) Car insurance: why women face £300 rise in premiums. *The Guardian* October 5th 2012. <<https://www.theguardian.com/money/2012/oct/05/car-insurance-women-rise-premiums-eu-ruling>> Last accessed July 18th 2017.

- Cudd, Ann. E. (2006) *Analyzing oppression*. Oxford: Oxford University Press.
- Darwall, Stephen L. (1977) Two kinds of respect. *Ethics* 88(1), 36-49.
- Eidelson, Benjamin (2015) *Discrimination and disrespect*. Oxford: Oxford University Press.
- Ezorsky, Gertrude (1991) *Racism and justice: the case for affirmative action*. Ithaca, NY: Cornell University Press.
- Fara, Delia Graff (2001) Phenomenal continua and the sorites. *Mind* 100(440), 905-935.
- Fiss, Owen M. (1976) Groups and the equal protection clause. *Philosophy & Public Affairs* 5(2), 107-177.
- Frye, Marilyn (1983) Oppression. In Terence Ball and Richard Dagger, eds. *Ideals and ideologies: a reader*, 8th edn., 375-383. Boston, MA: Pearson.
- Goldman, Alan H. (1979) *Justice and reverse discrimination*. Princeton, NJ: Princeton University Press.
- Graham, Peter A. (2010) In defense of objectivism about moral obligation. *Ethics* 121(1), 88-115.
- Harris, David A. (2003) *Profiles in injustice: why racial profiling cannot work*. New York, NY: The New Press.
- Hellman, Deborah (2014) Racial profiling and the meaning of racial categories. In Andrew I Cohen and Christopher Heath Wellman, eds. *Contemporary debates in applied ethics*, 2nd edn., 232-243. Malden, MA: Wiley-Blackwell.
- Jackson, Frank (1991) Decision-theoretic consequentialism and the nearest and dearest objection. *Ethics* 101(3), 461-482.
- Kagan, Shelly (2011) Do I make a difference? *Philosophy & Public Affairs* 39(2), 105-141.
- Keefe, Rosanna (2011) Phenomenal sorites paradoxes and looking the same. *Dialectica* 65(3), 327-344.
- Kennedy, Randall (1999) Suspect policy. *The New Republic* 221(11-12), 30-35.
- (2002) Blind spot. *The Atlantic Monthly* 289, 24.
- (2013) *For discrimination: race, affirmative action, and the law*. New York, NY: Pantheon.
- Kernohan, Andrew (1998) *Liberalism, equality, and cultural oppression*. Cambridge: Cambridge University Press.

- Khaitan, Tarunabh (2015) *A theory of discrimination law*. Oxford: Oxford University Press.
- Kutz, Christopher (2000) *Complicity: ethics and law for a collective age*. Cambridge: Cambridge University Press.
- Lang, Kevin and Lehmann, Jee-Yeon K. (2012) Racial discrimination in the labor market: theory and empirics. *Journal of Economic Literature* 50(4), 959–1006.
- Lever, Annabelle (2005) Why racial profiling is hard to justify: a response to Risse and Zeckhauser. *Philosophy & Public Affairs* 33(1), 94–110.
- Levin, Michael (1992) Responses to race differences in crime. *Journal of Social Philosophy* 23(1), 5–29.
- Levy, Neil and McKenna, Michael (2009) Recent work on free will and moral responsibility. *Philosophy Compass* 4(1), 96–133.
- Lippert-Rasmussen, Kasper (2006) Racial profiling versus community. *Journal of Applied Philosophy* 23(2), 191–205.
- (2007) Nothing personal: on statistical discrimination. *The Journal of Political Philosophy* 15(4), 385–403.
- (2014) *Born free and equal? A philosophical inquiry into the nature of discrimination*. Oxford: Oxford University Press.
- MacKinnon, Catharine A. (1979) *Sexual harassment of working women: a case of sex discrimination*. New Haven, CT: Yale University Press.
- Meehan, Albert J. and Ponder, Michael C. (2002) Race and place: the ecology of racial profiling African American motorists. *Justice Quarterly* 19(3), 399–430.
- Miller, David (1999) *Principles of social justice*. Cambridge, MA: Harvard University Press.
- Nagel, Thomas (1973) Equal treatment and compensatory discrimination. *Philosophy & Public Affairs* 2(4), 348–363.
- Nefsky, Julia (2011) Consequentialism and the problem of collective harm: a reply to Kagan. *Philosophy & Public Affairs* 39(4), 364–395.
- Parfit, Derek (1984) *Reasons and persons*. Oxford: Oxford University Press.
- Pojman, Louis P. (1998) The case against affirmative action. *International Journal of Applied Philosophy* 12(1), 97–115.

- Prichard, H. A. (1932) Duty and ignorance of fact. In H. A. Prichard *Moral writings*, 84–101. Oxford: Oxford University Press.
- Ramirez, Deborah A., Amy S. Farrell, and John McDevitt (2000) *A resource guide on racial profiling data collection systems: promising practices and lessons learned*. Washington, DC: US Department of Justice.
- Risse, Mathias (2007) Racial profiling: a reply to two critics. *Criminal Justice Ethics* 26(1), 4–19.
- Risse, Mathias and Zeckhauser, Richard (2004) Racial profiling. *Philosophy & Public Affairs* 32(2), 131–170.
- Ross, W. D. (1939) *Foundations of ethics*. Oxford: Oxford University Press.
- Schauer, Frederick (2003) *Profiles, probabilities, and stereotypes*. Cambridge, MA: Harvard University Press.
- Segall, Shlomi (2012) What’s so bad about discrimination? *Utilitas* 24(1), 82–100.
- Singer, Peter (1978) Is racial discrimination arbitrary? *Philosophia* 8(2–3), 185–203.
- Spiekermann, Kai (2014) Small impacts and imperceptible effects: causing harm with others. *Midwest Studies in Philosophy* 38(1), 75–90.
- Srinivasan, Amia (2015) Normativity without Cartesian privilege. *Philosophical Issues* 25(1), 273–299.
- Sunstein, Cass R. (1994) The anticaste principle. *Michigan Law Review* 92(8), 2410–2455.
- Thomas, George C. (2001) Blinded by the light: how to deter racial profiling – thinking about remedies. *Rutgers Race and the Law Review* 3(1), 39–59.
- Thomson, Judith Jarvis (1973) Preferential hiring. *Philosophy & Public Affairs* 2(4), 364–384.
- Voorhoeve, Alex and Binmore, Ken (2006) Transitivity, the sorites paradox, and the similarity-based decision-making. *Erkenntnis* 64(1), 101–114.
- Wasserman, David T. (1991) The morality of statistical proof and the risk of mistaken liability. *Cardozo Law Review* 13(2–3), 935–976.
- (1997) Diversity and stereotyping. *Philosophy and Public Policy Quarterly* 17(1–2), 32–36.
- (2011) Is racial profiling more benign in medicine than law enforcement? *Journal of Ethics* 15(1–2), 119–129.
- Woodard, Christopher (2008) *Reasons, patterns, and cooperation*. New York, NY: Routledge.

Young, Iris Marion (1990) *Justice and the politics of difference*. Princeton, NJ: Princeton University Press.

Zack, Naomi (2015) *White privilege and black rights: the injustice of U.S. police racial profiling and homicide*. Lanham, MD: Rowman & Littlefield.

Zimmerman, Michael J. (2006) Is moral obligation objective or subjective? *Utilitas* 18(4), 329-361.