

The Art of the Unseen – Three challenges for Racial Profiling

Frej Klem Thomsen

fkt@ruc.dk

Department of Philosophy & Science Studies

Roskilde University

Denmark

The Art of the Unseen – Three Challenges for Racial Profiling

“L'art de la police est de ne pas voir ce qu'il est inutile qu'elle voie.” (Napoleon Bonaparte, Letter to Minister of Police Joseph Fouché, May 24, 1800)

Abstract. This article analyses the moral status of racial profiling from a consequentialist perspective and argues that, contrary to what proponents of racial profiling might assume, there is a prima facie case against racial profiling on consequentialist grounds. To do so it establishes general definitions of police practices and profiling, sketches out the costs and benefits involved in racial profiling in particular and presents three challenges. The foundation challenge suggests that the shifting of burdens onto marginalized minorities may, even when profiling itself is justified, serve to prolong unjustified police practices. The valuation challenge argues that although both costs and benefits are difficult to establish, the benefits of racial profiling are afflicted with greater uncertainty than the costs, and must be comparatively discounted. Finally, the application challenge argues that using racial profiling in practice will be complicated by both cognitive and psychological biases, which together reduce the effectiveness of profiling while still incurring its costs. Jointly, it is concluded, these challenges establish a prima facie case against racial profiling, so that the real challenge consists in helping officers practice the art of the police and not see that which it is useless that they should see.

Keywords: Consequentialism, Crime, Ethics, Police, Profiling, Race

Introduction: Controversy and Intuitions

“Profiling” can be provisionally defined as “the application of statistical evidence concerning differences in propensity for crime in a police practice when deciding whom to target.” Racial profiling is thus the use of statistical evidence concerning racial differences in crime-rates in the attempt to increase the likelihood of apprehending offenders.¹ It is both in the public debate and the professional literature an enormously controversial topic, with strong opinions both for and against but not as much clarity as one might prefer about what *exactly* the pros and cons of racial profiling are. In the following, I aim to explore these in some detail and sketch what I take to be three serious challenges for the proponents of racial profiling.

Let me begin with some fundamental observations. Much of the controversy regarding racial profiling stems, I believe, from the presence of two conflicting intuitions about the imposition and distribution of costs on those subjected to police practices. Consider the following scenario, a slight variation on a set of actual cases described by David Harris (Harris 2002):

Chicago Customs. You are returning from a visit abroad and re-entering your country at the airport. While going through customs you have your bags hand searched by a customs agent, carefully looking through every single item in your suitcase, backpack and pockets. You are questioned extensively about your travelling origin, purpose for travelling, destination, occupation, home address, etc. At the end you have to repack your luggage yourself. You are then directed toward a closed room off to the side. You ask for an explanation, but receive none. When you complain, you are yelled at by one agent, and the other

¹ How best to understand race is a both complicated and controversial issue. (Cf. Angelo Corlett’s article in this issue of *The Journal of Ethics*) In the following discussions of race, I do not mean to imply that there is a genetic basis for distinguishing between biologically distinct races, only that race is clearly part of the social landscape as a trait which is present in the perceptions of people, whether by identification of or self-identification with a race. (Loury 2005) Furthermore, and perhaps more problematically, I shall use “race” as essentially synonymous with “ethnicity.” This tracks the way the focus in the United States context tends to be “racial profiling,” whereas in the European context it is “ethnic profiling,” although the two are relevantly identical phenomena. Although this conflation is somewhat less conceptually refined than I should prefer, I do not believe that any accompanying problems affect my central arguments.

remarks: “Oh no, not another one...” Once inside the closed, windowless room, you are placed hands against the wall, and searched thoroughly by the agents, who run their hands tightly over your entire body including your genitals. When nothing is found you are strip searched, and when that too reveals nothing you are body cavity searched. Finally, having turned up nothing, the agents decide to apply a “monitored bowel movement,” ordering you to imbibe a powerful laxative and wait for you to empty yourself into a conveniently placed non-flushable toilet. Only after several hours, when you have disclosed the contents of your digestive system to their content are you finally released. At no point do you receive an explanation for the treatment, and your explicit requests for legal aid and contact to your family, who expected you home a while ago, are consistently denied.

Most people would, I believe, be indignant and angry at being treated in the way described in *Chicago Customs*, and critical of a police practice which regularly and systematically subjected persons to such treatment. And rightly so. For although there are aggravating circumstances that could conceivably be remedied while maintaining the practice, such as the callousness of the agents handling your case, it seems clear that the bulk of your complaint concerns the discomfort, inconvenience and humiliation that you suffer as a consequence of the intensive scrutiny. This police practice imposes a heavy cost on its targets.

Now consider a different scenario, loosely based on data regarding stop-searches in London (London Metropolitan Police Authority 2010):

London Metropolitan. You live in an ethnically diverse neighbourhood with both Black, Asian and White occupants. Following a violent crime the police intensify their use of stop-searches, and one day on the street you are confronted by two officers who order you to submit to a search. They carefully search you by running their hands tightly against your entire body while you lean on a wall, legs spread and palms firmly on the bricks. A week later you are stopped again, and this time the officers escort you into an open gateway where, only partially shielded from public view, you are required to strip down to your underwear as they

search your clothing. At the end, one officer pulls your underwear away from your body, while briefly shining a flashlight into it and checking your genitals. Three days later you face your third search, although this time only of the “pat-down”-variety you first experienced. Discussing the frequency with which you have been searched with friends, you find deep disparities along ethnic lines, and when you look up the statistics it turns out that members of your ethnicity have at least a four times greater frequency of being searched than members of the dominant ethnicity.

Most of us would also, I think, intuitively find the police practice in *London Metropolitan* problematic. Surely, it cannot be right that the police discriminate against persons of a certain ethnicity to the extent that such persons are searched four (or more) times as frequently as others? Yet, and here is the rub, the intuitions I have attempted to entice in these two cases can, and often do, pull in opposite directions.

The imposition of such severe costs, in terms of e.g., inconvenience, discomfort and humiliation, as those described in *Chicago Customs* requires that the police have good reasons for imposing these costs. And such reasons it seems must consist in either a sufficiently great benefit in positively identifying a criminal – suppose that the police are searching for a smuggler carrying the final component necessary for a terrorist organization to arm and use a nuclear device – a sufficiently great likelihood that the persons subjected to the practice will be positively identified – suppose that the police know with complete certainty that someone matching a specific description will arrive with your flight and be carrying drugs, and that only you and one other person on your flight match the description, giving each of you a 0.5 probability of being the courier – or, ideally, a combination of the two.² In *Chicago Customs* it seems to me clear that selecting targets completely at random, each of whom would have an equally and (presumably) very small probability of being positively identified as a smuggler of any kind, would be clearly wrong since the burden imposed on them is so great. In fact, the *only* way that the police can justifiably impose burdens of such calibre, given that they cannot alter the benefits of detecting a smuggler, is to carefully select for

² I shall return to and expand on why I consider this to be the case in the discussion of costs, benefits and justifiability below.

scrutiny only those persons who have a much greater than average probability of being smugglers.³ Profiling is *necessary* to make such cost-severe police practices justifiable. Historically, what brought attention to the cases described by Harris was that they targeted black women, clearly engaging in both racial and gender profiling. Added to the wrong many felt was committed by discriminating in this way was the fact that hit rates for black women were lower than for comparative demographics, so that targeting this group seems to have been irrational in the narrow sense of *decreasing* the likelihood of successfully identifying drug smugglers. But the trigger for the outrage that followed is explained in terms of the fact that racial profiling was applied, whereas as I have stated above my initial feeling is that such practices would be unjustified no matter who they happened to, unless careful profiling meant that the practice had a high likelihood of apprehending criminals.

The pull of *London Metropolitan* is in the opposite direction. Although presumably any one of the searches described would have been uncomfortable the cost imposed is much less and therefore easier to justify than in *Chicago Customs*; it is their focused and cumulative effect that is deeply distressing. In the actual case of Metropolitan London, the monthly rates of stop-searches for the period of May 2009-May 2010 varied between 3.09 per thousand and 4.94 per thousand for whites, and between 12.89 per thousand and 18.38 per thousand for blacks. (London Metropolitan Police Authority 2010) Asians, fluctuating between 5.42 and 8.34 per thousand, are much closer to the rates for whites, although still far more targeted.⁴ What troubles us here, I think, is that the profiling targets particular ethnicities and we tend to assume that for police to discriminate along

³ This is perhaps putting it too strongly: “only” here is meant only to refer to potential justifications within the framework considered in this article, that is, consequentialist justifications. As per the note above, I expand on this in the pertinent sections below.

⁴ Figures from all of England and Wales suggest even stronger discrepancies between the various ethnic groups. Thus in 2008/09, which is the latest data available, the rate of stop-searches for whites was 18.6 per 1000, while that of blacks was a staggering 134.3 per 1000, a ratio of ca. 1:7. In both cases, these levels represent a steep increase in the rate of searches since 2004/05, but for whites the increase is 25.7% (up from 14.8 per 1000) while for blacks it is 55.4% (up from 86.4 per 1000). (Statistics on Race and the Criminal Justice System 2008/09 2010; cf. also Phillips and Bowling 2002) The 1984 PACE-regulation allows such stop-searches when there is “reasonable suspicion” that the person subjected to them carries contraband, e.g., stolen goods or illegal drugs. Note that the United Kingdom is, to my knowledge, the only European country that systematically collects data on the ethnicity of those persons who interact with the criminal justice system, and therefore the most readily accessible source for figures such as the one I cited. But although verifying this is thus difficult, I see no reason to believe that the U.K. is unique among its European neighbours in its profiling practices.

racial or ethnic lines is either intrinsically or instrumentally bad.⁵ It might be intrinsically bad if e.g., police are obligated by considerations of justice to set aside all differences of gender, race and religion as irrelevant, that is, if justice requires the execution of the law to be “colour-blind.” It might be instrumentally bad if, e.g., targeting vulnerable minorities will offend and stigmatize them, the imposition of costs upon them will further burden the already deprived or the cumulative effect of the burdens will cross a threshold beyond which the harm caused grows severe enough to matter morally in a different way. In any of these cases, the thrust of the intuition runs counter to the notion that profiling can be required for a police practice to be justified, in that it points out the ways that profiling can delegitimize an otherwise justified police-practice. Caught between the requirements of minimizing the imposition of costs on the innocent and of avoiding the unequal distribution of burdens onto racial minorities, where does this leave racial profiling?

The answer, I will suggest, depends on the balance of costs and benefits in a way not essentially different from how we ought to assess police practices more generally. However, contrary to what is frequently assumed, such consequentialist considerations do not support racial profiling, as a strong *prima facie* case can be made against it based on the relative certainty of the associated costs and the relative uncertainty of the associated benefits. Given this *prima facie* case against it the onus shifts onto proponents of racial profiling, but the same uncertainties which support the *prima facie* case make it unlikely that proponents will be capable of producing a solid counter to meet the *prima facie* challenge.

In the following, I first set out at some length the background for the discussion, including more stringent definitions of the central terms and a minimal account of the relevant costs and benefits and several versions of a principle of justification. Much of the discussion in this part of the article, the account of costs excluded, is of a general character, applying to police practices and profiling broadly. In the next section, I then suggest three challenges that racial profiling in

⁵ Strictly speaking, we cannot infer that such discrimination is going on from the mere correlation between race and likelihood of being stopped. It is possible that both race and criminality correlate with some other trait, such as e.g., socio-economic deprivation, which is the trait actually profiled for. This would make it a case of indirect rather than direct discrimination. Further statistics, however, seem to indicate that police employ more profiling, and more prejudiced profiling, the more at liberty they are to do so: the ratio of black-to-whites in stop-searches under section 60 legislation, which has less strenuous requirements than “reasonable suspicion”, increase to 16.9 to 1, while success-rates, i.e., stop-searches resulting in arrests, plummet from 12% to 4%. (Open Society Justice Initiative 2009) In the following I shall not consider such complications, but focus the discussion on those cases where it is actually race or ethnicity that is being profiled.

particular encounters concerning its valuation, application and foundation. In conclusion, I argue that given these challenges and the difficulty of meeting them, racial profiling must be considered unlikely to be justifiable in realistic scenarios.

Policing and Profiling

Having briefly explored the background intuitions let me set out in more detail the central concepts of the discussion, specifically what it means to profile and how it relates to police-work more broadly.

By a “police practice,” I understand roughly a limited set of actions performed routinely by police officers for a specific and pre-defined purpose. Patrolling a certain neighbourhood and stopping certain people to question and potentially search them to find drugs would be an obvious example. It consists of a set number of reasonably well-defined and routine actions which can be performed easily in much the same way by any police-officer ordered to perform the practice. But going to the house of a witness to question her, pulling over and testing a driver for inebriation or performing an arrest could also be understood to be practices. While leaving this definition deliberately broad, however, I want to emphasize an important delimitation that it is worth bearing in mind. Unlike the focus of the present article much if not most of the actual work done by police officers is not concerned with the apprehension of criminals. “Police activities,” although different from one country to another, typically involve the ubiquitous paperwork that accompanies modern institutions, including such service-functions as the issuing of driver’s licences, criminal records and passports, as well as e.g. traffic regulation, maintaining public order, community work, increasing perceived safety through visibility, accident control, and the many miscellaneous forms of “helping out” from returning lost children to their parents to making rowdy teenagers turn down the volume of the stereo at house-parties. I set all these aside for present purposes, so that in the following when I refer to police practices, I shall in fact mean only those practices whose (immediate) purpose is to apprehend criminals. These include everything from the narrowly focused investigation of a serious crime to very broadly focused street-patrols intended among other purposes to allow officers to keep an eye out for any crime they happen upon. Of particular interest in the present context is the middle-ground between these two, occupied by the kinds of relatively focused screening exemplified by random alcohol-tests administered to drivers on specific days, in my native Denmark typically Fridays in December when there is

traditionally a high frequency of drunks returning from company Christmas-lunches. It is in such relatively focused practices that profiling is most obvious and easiest to apply.

By “profiling” we normally understand the use by police of statistical evidence as a tool for identifying and apprehending criminals. The fact that it is merely one tool among others is important, because as Mathias Risse and Richard Zeckhauser have pointed out there is a tendency to mistakenly equate racial profiling with practices where race is the only or determinate trait governing which persons are subjected to the practice. (Risse and Zeckhauser 2004) Thus with a slight broadening and rephrasing of a definition suggested by Risse & Zeckhauser we might initially say that profiling consists of “any police-initiated action that relies on statistical evidence and *not merely* on the behaviour of an individual.”⁶ (cf. Risse and Zeckhauser 2004)

I believe, however, that some further modifications are in order. First, note that the notion of statistical evidence at stake here is not rigorous. Officers need not know the exact ratio between male and female crime-rates to profile on the basis of gender; if they give special attention to male potential suspects based on their justified and true, if inexact, belief that men are vastly more likely than women to commit certain crimes, then this seems to me to qualify as profiling. It will be less efficient profiling the less precise and accurate their beliefs are about the actual ratios, but profiling nonetheless. Second, it is not obvious why we would want to draw a line between behavioural and other characteristics as Risse & Zeckhauser seem to suggest.⁷ Thus, if customs agents give particular scrutiny to plane-passengers who purchased their ticket in cash, wear sunglasses while checking in and travel short-term (all behaviours) based on statistical evidence that such persons have a greater than average probability of being drug-smugglers, this seems as clear a case of profiling as the special scrutiny given to men above. What might mislead us here is the intuition that the kind of evidence obtained from immediately observable behaviour – from appearing nervous to brandishing a gun – is in some relevant sense qualitatively different from the kinds of evidence which involves inferences from general facts about the world, such as the

⁶ Risse and Zeckhauser’s original formulation is that racial profiling consists in “any police-initiated action that relies on the race, ethnicity, or national origin and not merely on the behavior of an individual.”

⁷ Arguably, this is not the best reading of Risse & Zeckhauser’s definition, which is intended to encompass only racial profiling in contrast to, perhaps, other forms of profiling. But if so, a more natural phrasing would be that racial profiling is any police action that relies on statistical evidence about correlations between race and crime *in addition to* evidence about correlations between other traits, behavioural and non-behavioural, and crime. This is close to the definition that I shall suggest below.

correlation between purchasing a plane ticket in cash and smuggling drugs. This would be a mistake, because strictly speaking *no* evidence can avoid such inferences. Even taking nervous or threatening behaviour as evidence relies not only on inferences about what constitutes nervous or threatening behaviour – the relation between, e.g., facial expressions and body language on the one hand and states of mind on the other, and the notorious locus of intercultural misunderstandings – but also on the assumed (and presumably correct) correlation between such forms of behaviour and criminality. Nor can we avoid such deductive reasoning. We all necessarily and continuously form expectations about the traits other persons possess based on our observance of which traits they have demonstrated so far and our experiences of how such traits tend to correlate with others. In the words of Frederick Schauer: “It is simply how we think.” (Schauer 2003: 75) It is not a crime to appear nervous, but police officers may have good reason to be more interested in persons who appear nervous than in those who do not if such behaviour correlates with criminality.⁸ And exactly the same reason obtains in the cases of “non-behavioural” and statistical evidence.

A parallel distinction might be thought to more successfully track our intuitions: that of the difference between reactive and proactive policing, and as an extension thereof, between allowing racial and ethnic characteristics to be used to apprehend criminals when based on evidence about a particular criminal and allowing its use based on statistical evidence about criminals in general. By “proactive policing” I understand any activity by the police which is initiated by the police, not as a response to knowledge that a specific crime has been committed, but on the suspicion that a crime *may* have been (or is in the progress of being) committed. “Reactive policing,” by contrast, concerns situations where police respond to information that a crime has been committed.⁹

“Proactive police profiling,” then, means the employment of statistical evidence in searching for and apprehending randomly selected offenders. Stop-searches, which are both the most controversial instances of profiling in public debates and at the core of the discussion in

⁸ It may not on such a broad a description. Probably, there are particular kinds of “appearing nervous” which warrant attention and others that do not. Nor am I suggesting that “appearing nervous” is likely to be sufficient in and of itself to raise probabilities to a level where subjecting such persons to a police practice is justified. I mean only that they can count in favour of doing so. I elaborate on these points in connection with the discussion of proxy sets and justifiability below.

⁹ Here, as noted above, I limit myself to those police activities and practices whose purpose is the apprehension of criminals. Obviously, much reactive and in particular proactive policing will have different purposes, e.g., by striving to reduce crime-levels through other means than the apprehension of criminals.

academic debates, are of this kind. “Reactive police profiling” means the employment of statistical generalizations in searching for and apprehending one specific offender. An obvious example of this is the kind of profiling famously employed by behavioural criminologists to identify and apprehend serial-killers.

More importantly, proactive and reactive policing may seem to be morally different, even in cases where they both involve giving special attention to persons because of their race. Thus, some would want to argue that doing so in reactive policing is not morally wrong in the way doing so in proactive policing is. As David Harris explains it: “It does make sense to use racial or ethnic characteristics in enforcement, but only in one context: *cases in which race or ethnic characteristics describe actual suspects*. [...] In this situation, a description of the suspect’s skin color serves not as a predictor of criminality, but as an identifying physical attribute that can be used, in conjunction with others, to determine whether a person observed might be someone wanted by the police as a suspect in an actual crime. [...] What must *not* be allowed is using race or ethnic appearance, alone or in combination with other factors, to stop a particular person based on a *prediction* that he or she is more likely to be involved in crime.” (Harris 2002: 152; cf. also Wasserman 1996; Kennedy 1997)

The problem with Harris’ argument is that in both reactive and proactive policing it is exactly when and if a person’s traits, such as race, serve as a predictor of criminality that the traits become interesting to police officers. Consider how we might describe the epistemic situation of the agent who initiates the police-action. It is tempting to focus on what the evidence reveals about the offender, and to distinguish here between general evidence about groups of offenders and individualized evidence about a particular crime. But this is not the locus of the decision to subject persons to police-action. Whether given statistical or individualized evidence, police must first estimate how this affects the likelihood that an/the offender has a given trait (or set of traits), but second, and crucially, consider what the probability is that *this particular person* is an/the offender, *given* the likelihood now established that an/the offender possesses certain traits. While police-work thus necessarily involves both forms of inference, it is the later that is at stake in the decision whether to subject a person to a practice. If an eye-witness to a crime has described the offender as being of race R, then being of race R raises the probability that one is the offender, in just the same way that statistical evidence about, e.g., how gender correlates with crime does. This is what it

means “to determine whether a person observed might be someone wanted by the police...” The relevant difference does not concern prediction. (cf. Harcourt 2004)

A better explanation of the difference is perhaps the different proxies they involve. Thus: “It is not *racial* profiling for an officer to question, stop, search, arrest, or otherwise investigate a person because his race or ethnicity matches information about a perpetrator of a specific crime that the officer is investigating. That use of race – which usually occurs when there is a racially specific description of the criminal – does not entail a global judgement about a racial or ethnic group as a whole.” (Gross and Livingston 2002: 1415, emphasis provided) This does not mean, however, that it involves *no* “global judgments about a group as a whole,” or is not profiling. What it is, rather, is what we might call “testimonial-profiling,” in that officers are moving from, e.g., eye-witness testimony via an inference involving the general validity of such testimony to the probability that someone who matches the testimony is the offender.

This is still an important difference, because different forms of profiling, i.e., the application of different proxy-trait correlations, will look very different morally. Some will perform better, that is, be easier to apply, more accurate, less vulnerable to bias, and some worse, that is, be more costly to use, less efficient, etc. We will return to this shortly, in the next section on costs and benefits.

A final consideration is worth introducing at this point: the difference between single proxies and proxy-sets. I have been concerned in the previous primarily with the use of single proxies, and in particular with the proxy of race, but in fact we normally never correlate between just one trait such as race and expected behaviour, but adjust for, e.g., gender, age, dress-code, facial expression, body-language and social context, even to form such superficial impressions as those available to police in stop-search practices. As Laurence Thomas observes: “With typical social monitoring, it is quite normal to mark the difference between a white male stranger in a tweed coat and a white male stranger in gang-like garb, and to suppose that an isolated encounter with the former is less likely to be hostile than an isolated encounter with the latter. There is absolutely nothing about black males to suggest a different judgment is warranted between either a black male in a tweed coat and a black male in gang-like garb or, for that matter, a black male in a tweed coat and a white male in gang-like garb.” (Thomas 1992: 32-33) In short, clothing-style – as well as quite a few other traits – is normally a much better proxy than race for predicting behaviour, although contrary to what Thomas suggests, race might still enter the picture *in combination* with

any other relevant correlations in situations where race correlates with behaviours, e.g., crime.¹⁰ When considering whether or not to subject a person to a practice, whether it be conducting a random stop-search or questioning suspects of a crime, police officers will apply not just single proxies but a proxy-set to estimate the likelihood that said person has committed a crime. Some proxy-traits may negatively correlate, making it *less* likely that the person has committed a crime, and balance other proxy-traits that positively correlate. Whether this leaves e.g. the nervous, old, poor, white woman or the calm, young, wealthy, black man as the statistically most probable offender will be a difficult question to answer, but it is not one police officers can avoid and the only way of dealing with it is to weigh all the relevant factors.

Thus, a proxy-set for a police-practice is a set of proxy-traits which jointly establish the probability of any given person being an identifiable offender,¹¹ so that for any given proxy-set:

Proxy Set S. S is the set of proxy-traits (T, T₁, T₂...) that jointly establish probability P_S that a randomly selected person A with trait(s) T (and/or T₁, T₂...) will be identified as having committed a criminal act if subjected to police-practice PP.

Naturally, no proxy-set is better than the correlations on which it rests. These need to be non-spurious, and as accurately and precisely specified as possible. This is important, because there are literally an infinite number of traits that could be considered to be proxy-traits for criminality. The art of the police consists, as the quote which inspired the title of this article suggests, in ignoring the infinity of traits that have no bearing on the issue and picking out those traits that do.

¹⁰ Thus, Thomas concludes: “Regardless of the percentage of black hoodlums, it is rational to extend the statistic to all black youths only if black hoodlums cannot be reasonably and safely distinguished from other black youths, given a modicum of social monitoring skills and prevailing norms of self-presentational behavior.” (Thomas 1992: 34; cf. also Cox 1993) But, we might well insist, even in those situations where they can be distinguished in such a way a combination of clothing-style, race and any other relevant traits would, in the relevant circumstances, be a better set of proxies than clothing-style by itself.

¹¹ Note that strictly speaking we need to distinguish sharply between the proxy-set which assigns probabilities that a person is an offender and the proxy-set which assigns probabilities that offenders will be identified by the practice in question. Although the two sets will likely contain mostly the same proxies they will assign different probabilities, given the inevitable imperfections of any practice in apprehending offenders subjected to it. In practice we are concerned primarily with the latter, and I shall restrict myself to this in the following, setting the differences between the two and the complexities introduced by distinguishing them aside.

Let me sum up the main points that I have attempted to establish in this section. Police use profiling in their various practices intended to apprehend offenders. By “profiling” we might initially understand the use of statistical evidence to better target offenders, but I have argued that there is no sharp dividing line between the way of targeting offenders which we associate with paradigm cases of profiling and other ways of targeting offenders; they simply apply different proxies in different ways. I have further argued that this also holds true across the division between reactive and proactive policing, that is, police actions that respond to known crimes that have been committed and those that attempt to detect or prevent unknown crimes, typically those in progress. Although it is tempting to describe these as morally different, and to allow otherwise impermissible characteristics, such as race, to enter police-work in the former case, the real difference concerns the proxy-traits at stake rather than the difference in kinds of police-work. Finally, I have suggested that profiling rarely, if ever, concerns only one proxy-trait. Instead, it will normally involve a set of traits which jointly provide police with (an estimate of) the probability that a person has committed a crime and can be apprehended if subjected to a given practice.

Thus, I would suggest, police profile for proxy trait T on practice PP *iff*:

1. Police believe T is part of proxy-set S (T, T1, T2 ...), the members of which correlate with a relevant target-trait (criminality);
2. Police use S to target persons for PP, by basing their selection of which persons to subject to PP on the probabilities assigned by S.

Given this understanding of profiling as a broad and inevitable phenomenon, the question that I shall be concerned with in the remainder of this article is not “to profile or not to profile,” but how.

The Benefits and Costs of Policing and Profiling

The strongest arguments for the use of racial profiling, it seems to me, must be consequentialist. Racial profiling, proponents could argue, help police to better target offenders when and if race correlates with crime, allowing police to catch more criminals. This is a beneficial consequence, so the argument will go, which justifies the imposition of an otherwise distasteful difference in treatment along racial lines. If not exactly common in the philosophical literature, this

is certainly a line that is frequently suggested in public debates. However, as I hope to show in the following, the situation when viewed within a consequentialist framework is both considerably more complex and considerably less propitious for racial profiling. To demonstrate this I shall first explore the structure of costs and benefits in policing and the way that profiling fits into this. Secondly, I shall outline three challenges that I believe racial profiling specifically faces.

This argument may be of interest, I hope, to non-consequentialists as well for two reasons. First, most forms of non-consequentialism will want to allow consequences to matter to some extent, even if the weighing of consequences is held to not constitute the full moral picture. Secondly, and perhaps more importantly, if as it seems to me the strongest arguments for racial profiling are consequentialist, and racial profiling fails on consequentialist terms, then there is little left to be said in its defence. If it cannot be defended on these grounds, then proponents will be very hard pressed indeed to argue their case.

Consider first the standard costs and benefits we might attribute to a police practice intended to apprehend criminals. These will vary, of course, depending on the ultimate value theory applied by the specific variant of consequentialism at stake, but most forms of consequentialism will converge on holding at least the following factors to have instrumental value, irrespective of the intrinsic value(s) which they ultimately serve:

Costs and benefits of criminal apprehending police practices

Resources	Deterrence
Inconvenience	
Punishment	Incapacitation

“Resources” here means the expenditure of these in terms of human-hours of police-work and the economic costs to society of training, equipping and paying police officers to invest these hours. These are resources that could otherwise have been invested differently by society in various beneficial ways. Allocating them to police practices thus counts as a cost of these forms of policing. By “inconvenience” I mean to cover those costs that accrue to the persons subjected to the police practice, in terms of the time they are required to set aside for it, the discomfort, pain or humiliation

they may and will often suffer as well as any direct costs such as damaged clothing, vehicles or even personal injuries. Most persons, irrespective of issues of discrimination, will find, e.g., being stopped, questioned and physically searched by the police to involve such costs, even if only moderately so. The *Chicago Customs* case described in the introduction involves unusually severe inconvenience costs, but I do intend for this meek-sounding category to cover also the serious pain and humiliation that persons would likely suffer in that case. This cost, it is worth noting, applies to all persons subjected to the practice, both those found to be offenders and those who are not. Finally, “punishment” covers those costs imposed eventually by the criminal justice system on those persons found to be offenders by the practice. Given that they would not have had these costs imposed upon them had they not been subjected to the practice (at least, not on this occasion), the harm that criminal justice routinely inflicts on convicted criminals in the form of e.g. incarceration must be counted as part of the cost of the practice.¹²

On the benefits side, the two benefits that I have included are both traditionally taken to be valuable because they contribute to decreasing crime, and thus preventing the harms caused by crimes. “Deterrence” is thus the prevention of crime by increasing the threat to potential offenders to a level where they are dissuaded from committing a crime that would otherwise have occurred. “Incapacitation” is the prevention of crime by specifically stopping a potential offender from carrying out a crime, and ought therefore to cover both crimes that are detected while in progress by the offender being subjected to the practice and any crimes that the offender would have committed had she not been subjected to the practice but cannot as a result of being identified, i.e., because she is convicted and incarcerated.

Consider now what the costs and benefits of a given practice will be. Any realistic practice will only target a small minority of the population. Call these the positives, and the

¹² Here, and elsewhere in much of the following, my analysis potentially faces the challenge of double-counting, i.e., if the punishment of criminals is held to be a cost of the police practice which apprehends them, it might be argued, it cannot at the same time count as a cost of other parts of the criminal justice system where we would traditionally like to include it, such as, e.g., the courts which hands out the sentence or the criminal law which mandates it. While I recognize this as an important issue, the limitations of this article preclude dealing with it and the underlying question of the relation between causality and cost-attribution. More importantly, as noted above, the sketch I am providing of the costs and benefits of police work serves first and foremost to illustrate the costs and benefits of profiling, and to set the stage for the challenges I shall present. These central points are not substantively affected by whatever view on the issue of double-counting one adopts. As such, I leave it open to the reader to discount the cost of punishment (and the benefits to which we turn below) by the requisite amount, if she feels that this is necessary to avoid double-counting.

remainder the negatives. Further, for any realistic practice a group of the persons subjected to the practice will be innocent, indeed typically the majority will not be found to have committed a crime. Call these “false positives,” and those identified as offenders true positives. Some of the costs and benefits will apply to false positives, while other costs and benefits will apply to true positives.¹³ Notably, each instance of a false positive will incur full resource and inconvenience costs, but no punishment costs, and provide (presumably) partial deterrence but no incapacitation benefits. The benefit of partial deterrence can be presumed because even an instance of a police practice that fails to apprehend an offender is a public event which sends a signal to potential offenders that there is a risk of discovery and punishment. True positives on the other hand incur full costs and provide full benefits as per above. Finally, some costs and benefits will vary with the person upon which the practice is imposed. An ailing elderly person may suffer inconvenience more severely from being subjected to a police practice than a healthy adult subjected to the same practice. In the following I shall consider all costs and benefits to be averaged for the practice as a whole, so that each instance is considered to produce the same benefits and incur the same costs.¹⁴

Given the disparity between the costs and benefits provided by true and false positives, the ratio between them is clearly crucial to the costs and benefits of a police practice. Generally speaking, we assume that most, if not all, police practices rely on the positive cost-benefit ratio of their true positives to outweigh the negative cost-benefit ratio of their false positives. That is, although each instance of subjecting an innocent person to the police practice exacts a greater cost than the benefit it produces, the occasional instance of subjecting an offender to the practice

¹³ The costs and benefits of negatives are incorporated in the account of costs and benefits of positives, in that the benefit of, e.g., deterrence consists in the reduction of a cost imposed by the total population. These need therefore not be counted separately.

¹⁴ We need to carefully distinguish also between the costs that properly accrue to policing and profiling, and those that are often associated with but strictly speaking independent from policing and profiling per se. As Risse and Zeckhauser point out some of the wrongs often associated with racial profiling, such as the abuse of police powers to intimidate or harass civilians, are wrong independently of whether or not they occur in connection with profiling and even of whether or not they happen to be motivated by the same reasons which motivate profiling (such as racist beliefs and attitudes). What we are concerned with is whether and, if so, when racial profiling is wrong, which means whether it would also be wrong in situations where it did not coincide with other wrongs. Or, if we want to put it that way, whether racial profiling in and of itself constitutes a form of abuse by the police. (Risse and Zeckhauser 2004; cf. also Lippert-Rasmussen 2006; Lippert-Rasmussen 2007) However, while in the following I set aside costs arising from such phenomena as racism, police brutality and the abuse of powers, I do think an argument can be made for including such abuses as have been specifically enabled by a practice or a form of profiling. That is, if a practice or form of profiling adds to the amount of abuse, say by increasing the contact between racist officers and minority persons that they would not otherwise have encountered, then the resulting abuses ought to have some weight as costs of the practice.

produces a sufficiently great benefit, even considering its added punishment costs, that the practice in total is morally justifiable. This is the intuition driving the *Chicago Customs* case, where presumably the cost imposed on you by being searched as an innocent is much greater than any benefits produced, hence the only way to justify the practice as a whole is if there is a relatively high probability that people like you are offenders. It is not, however, a necessary truth that the cost-benefit ratio between true and false positives will take this form. We can imagine police practices, e.g., where even false positives have positive cost-benefit ratios, so that each time anybody, innocent or offender, is subjected to the practice, the total effect will be beneficial, or even practices where, perversely, false positives have positive cost-benefit ratios, but true positives have negative cost-benefit ratios, perhaps because of draconic criminal laws that lead to high partial deterrence benefits but extraordinary punishment costs, so that the practice relies on the instances of subjecting innocents to it to balance the instances of identifying offenders.

We can also, of course, imagine practices that fail to do more good than harm, if, e.g., both true and false positives have greater costs than benefits, or if the costs in one case are severe enough that the benefits of the other fails to outweigh them. Such scenarios need not be implausible.¹⁵ Indeed, if the success-rate is as low as in a recent Danish case of stop-searches, where a 4 week intensive stop-search initiative directed against possessing and carrying weapons resulted in 1 person being charged after 610 searches, we might well dismiss stop-searches irrespective of any issues raised by profiling. (Holst 2009) Nor is there reason to think that this low rate of success is extraordinary. Over a 14 month period, Danish police performed 17.977 searches, finding 300 weapons, a hit rate of 1.6%. (Lindqvist 2010) Similarly, in Amsterdam, the police spent 11.687 hours, conducted 32.332 searches and found 702 “weapons” (ca. 1 in 46) of which 15 were firearms (ca. 1 in 2155).¹⁶ (Open Society Justice Initiative 2009) I shall assume in the following, however, that we are discussing only practices that follow the first pattern identified, that is, practices where the positive cost-benefit ratio of true positives outweigh the negative cost-benefit ratio of false positives.

¹⁵ See Jesper Ryberg’s article in this issue of *The Journal of Ethics* for an argument that this is likely to be the case in realistic examples of racial profiling.

¹⁶ The figures for the Netherlands as a whole are no less worrying: “A 2005 Dutch study of the efficiency of preventive searches for weapons in eight cities over a two-year period found that the searches disproportionately targeted minorities and that the hit rate was only 2.5 percent [...] this figure was inflated by the inclusion of items such as penknives.” (Open Society Justice Initiative 2009)

In any of these cost-benefit distributions, the ratio between true and false positives is provided by the probability of the practice successfully identifying a randomly selected person as an offender given the proxy-set used to select targets for the practice. Further, if we set as the minimal requirement for a practice that it must do more good than harm,¹⁷ we can summarize the costs, benefits and justification in the following terms:

Costs, benefits and minimal justification of police practices:

1. Police practice PP has probability P_I of identifying a randomly selected person A as having committed a criminal act, i.e. $P_I = P(A_{\text{true}} | \text{PP})$;
2. True positives produce benefits B_{true} and costs C_{true} ;
3. False positives produce benefits B_{false} and costs C_{false} ;

PP is *minimally justified* iff. The expected benefits of PP outweigh the costs, i.e. $(P_I \cdot B_{\text{true}} + P_I \cdot C_{\text{true}}) + ((1 - P_I) \cdot B_{\text{false}} + (1 - P_I) \cdot C_{\text{false}}) \geq 0$.

The role of profiling in this should be clear. Since the ratio of true positives to false positives is determined by the probability of PP identifying offenders, and since this probability is a function of the proxy-set applied to target offenders, the profiling employed is crucial to the justifiability of a practice. The potential benefits of profiling are therefore reasonably obvious and uncontroversial: by applying non-spurious proxies to refine selection one will increase the benefits derived from apprehending an offender or decrease the costs in terms of police resources and inconvenience.¹⁸

Consider an example:

Profiling 101. Police practice PP₁ subjects 101 persons to the practice every day. These 101 persons are selected from a much larger group of

¹⁷ I shall return to what it means for a practice to be maximally justified, i.e., to meet the consequentialist requirement of being not only beneficial, that is to have better consequences than inaction, but optimal, that is at least as good as any available alternative action. Cf. the section on the valuation challenge, below.

¹⁸ Frequently, the first type of benefit, the increased apprehension of offenders, is granted the most attention. But the second type is important because it means that profiling may be justified, even in situations where apprehending offenders is not desirable. Suppose for instance that the effects on society-wide crime levels from apprehending more offenders are negligible, but that the law imposes draconic punishments on apprehended offenders. In that situation, it may not be desirable to increase the number of apprehended offenders, and so profiling cannot be justified in this respect. But it may still be the case that it increases police efficiency, so that it would be possible to apprehend the same number of offenders with fewer police resources, freeing those resources for other potentially beneficial types of police work, or even for altogether different policy initiatives.

10.100 potential subjects. Further, in each instance of subjecting a person to the practice police select the person from a much smaller group of presently available subjects, say 20, who they, based on proxy-set S_1 , consider to be the most likely offender. These have, as it turns out, an average probability of 0.05 of being an offender, meaning that roughly five offenders are caught per day, while 96 other subjects are false positives. Now we add non-spurious proxy T to S_1 , but otherwise leave the practice unchanged. S_{1+T} will allow police to be somewhat better, depending on the strength of T , at selecting the one person among the 20 with the greatest probability of being an offender. Those persons in the 20 who are T -persons will have their probabilities increased somewhat, to the extent that if the most probable offender on the basis of S_1 was a non- T -person, then a different person may now be the most probable offender on the basis of S_{1+T} . The difference between profiling and not profiling for T is the number of times that this happens. Every one of these actions has a slightly higher probability of apprehending an offender than the practice did just using S_1 . As it turns out, the persons selected by S_{1+T} have an average probability of 0.075 of being offenders. This means, if the number of actions remains unchanged, ca. 2.5 more true positives per day presumably increasing overall benefits, or, if the number of apprehensions is kept constant, a decrease of ca. 33 actions daily presumably decreasing overall costs.

Profiling, in short, increases the efficiency of a practice by better targeting, decreasing the number of false positives per action and increasing the number of true positives per action.

Note the importance of proxy-sets here. If we disregard the existing set when adding a proxy-trait, or if we imagine a situation where the set has only one member, then we would get a situation where any person with the trait is more likely to be an offender than any person without the trait: “Say eyewitness testimony suggests that there is a 60 percent chance that a crime was committed by an African American man, and African American males make up 25 percent of the population; one should then inspect *only* African American males, and *mutatis mutandis* for other scenarios. The reason is that an African American male is $2.4 = 60\%/25\%$ times as likely to be

guilty as a person selected at random.”¹⁹ (Risse and Zeckhauser 2004: 140) Although Risse & Zeckhauser point to several qualifications for this view, it still seems to me to inaccurately characterise the demands of efficiency, at least in any realistic scenario where police will have more proxies to go by than gender and race. Rather, the practice of inspecting will involve certain costs and benefits. For some P_I the expected costs and benefits of subjecting a person to inspection will balance. Call this threshold-probability N . Whenever police have the opportunity of subjecting a person with $P_S > N$ to PP the expected consequences are therefore beneficial, and as above, I shall call it minimally justified for a police agent to subject a person to police practice PP when that person’s probability of being an offender is equal to or greater than N . Conversely, police ought to refrain from subjecting persons with $P_S < N$ to PP; in these cases the expected costs will outweigh benefits. In summary:

Minimally justified police action. It is *minimally justified* for a police agent to subject person A to police practice PP iff $P_S \geq N$, where $N = P_I$ such that $P_I \cdot B_{true} + (1 - P_I) \cdot B_{false} = P_I \cdot C_{true} + (1 - P_I) \cdot C_{false}$.

Consider how this will affect police practices, in a variation of the example above:

Profiling 101a. Suppose that a neighbouring police department still using proxy-set S_1 reviews police practice PP_1 , and decides that $N = 0.1$. The existing practice thus does more harm than good and would continue to do so even with the refined selection of S_{1+T} – the bar needs to be set higher. As a result, police continue to apply S_1 but only subject persons where $P_S \geq 0.1$ to PP_1 . Suppose also that, unlike above, they never subject only the most probable in a group, but *all* persons encountered where $P_S \geq 0.1$. Their selection is still substantially

¹⁹ Our intuitions might steer us wrong here: if it is the case, e.g., that non-T-persons have 0.01 probability of being offenders and T-persons have a 0.04 probability then it might seem that an appropriate rate of stop and search is 1:4. But this is not true: in every single case where the police can decide whether to stop either a T-person or a non-T-person they are choosing sub-optimally if they choose to stop a non-T-person. An easily graspable comparison is with lottery-tickets, where we would not want to buy tickets with a 1% chance of winning if there are other tickets with a 4% chance of winning. Assuming that the tickets are identical in other respects, such as price and prize-size, the obvious choice is not to buy tickets in a ratio of 4:1, but to buy as many tickets with a 4% chance of winning as possible and to only start buying the 1% tickets once all of the 4% tickets have been sold (although strictly speaking, it is never rational to buy lottery tickets, because no lottery offers a better than even return for money).

narrower than before; they now subject only 40 persons daily on average, as these are the number of suitable targets locatable with same expenditure of resources.²⁰ Suppose that the resulting average probability is 0.125, resulting in roughly the same number of apprehensions (five daily), but a much lower number of false positives. The latter are the reason that PP is now overall beneficial. At this point T is introduced, leading to S_{I+T} being used to select targets. What happens? Much as above, some T-persons will see their probabilities shift above 0.1, while some non-T-persons may see theirs shift downwards below 0.1 (assuming negative correlation). The result? Some T-persons who ought to be subjected to PP, but would not have been under S_I , will be targeted using S_{I+T} , while some non-T-persons who ought not to be subjected to PP, but would have been using S_I , will not be targeted under S_{I+T} . As such the composition of the group targeted will change, and average probability of apprehension for PP will increase ($P_I < P_{I+T}$) with accompanying benefits.

Returning to Risse & Zeckhauser's example, if police use a proxy-set, rather than the single proxy provided by the evidence described, there may be whites and women with higher probabilities of being the offender than some black men, just as there may be blacks and men who do not have a sufficiently high probability of being the offender, in spite of the strong correlation established by the eye-witness testimony, to be legitimate targets for inspection. Imagine for instance that police know of one young, white woman who is a habitual committer of the crime in question, lives in the vicinity of the crime and has an appearance that could be mistaken for a young black man in conditions of confusion and stress, as well as of an elderly, law-abiding black man who was plausibly not in the vicinity of the crime at the time in question. Is it unrealistic to suppose that the first might have a high enough probability of being the offender to be worth inspecting, while the second does not? Possession of a trait that correlates with criminality makes it more likely that

²⁰ Note that the police department ought to increase resources until all persons with $P_S > N$ are subjected to PP, but doing so complicates matters because increasing costs in this way will raise N . There will be some point of optimal efficiency balancing these two conditions, but while finding it is a point worth pursuing in practice it is not relevant to the argument I am pursuing here.

one's probability of being an offender has crossed the threshold beyond which police ought to subject one to a practice, but it does not in and of itself guarantee it, no matter how strong the correlation.

What then of the costs of profiling? The costs most often associated with profiling concern its unintended side-effects, which I think can roughly speaking be said to fall in 4 groups: anti-deterrence, alienation, stigmatization and promoting structural inequality. Let me briefly sketch each of these in turn.

Anti-deterrent effects emerge from the fact that focusing police resources on one group will decrease the amount of resources available for actions against other groups. In the most extreme case, a police policy that targeted *exclusively* a specific social group which had been determined to be the most likely offenders for a type of crime would in effect remove the deterrent effect upon all persons not members of this social group, by guaranteeing them that they would not be subject to police scrutiny. (cf. Risse and Zeckhauser 2004; McGary 1996) In reality this cost is compounded, as Bernard Harcourt shows, by the way different groups may respond differently to marginal changes in the deterrent threat. To briefly restate his point, we can imagine a city with a deprived minority who are more prone to criminality but also highly resistant to variations in deterrence-incentives, perhaps because they lack adequate alternatives to the criminal career. What this means is that increased profiling of this group, though more efficient in terms of successfully targeting offenders, will produce a disproportionately small deterrent effect, easily offset by the rise in offences among the majority-group resulting from the decreased deterrence they experience. (Harcourt 2004)

Alienation effects stem from the fact that, irrespective of whether profiling is rational or not, there will be situations where those profiled experience the increased attention as an expression of distrust, disrespect and/or discrimination. This can have a number of implications, from less cooperation with the police force, to increased criminality because the legal system, its norms and its enforcers are no longer viewed as legitimate. At the extreme end, one frequently cited problem with profiling young Muslim men for terrorism is the risk that this will lead some of them to sympathise with the terrorism that the profiling is intended to prevent. While it is unlikely that profiling will in and of itself make a person support, much less actively partake in terrorism, the perceived injustice of being profiled could in some cases be the nudge that pushes a wavering potential criminal, terrorist or otherwise, to going through with the deed.

Stigmatization effects stem from the impression of the profiled group created by profiling, both on those persons subjected to profiling and on others (police officers, non-profiled members of the group, non-members of the group). While a profiling practice ought to rely on rational attitudes regarding the profiled persons, i.e., by applying statistically correct correlations, it is both possible and likely that profiling can engender irrational attitudes about the target group. Thus, as Arthur Applbaum notes, it is all too easy to slip from a rational correlation to over-generalizing, as well as to associating correlations with blameworthiness, and so for police officers: “As the proportion of true positives picked out by a strategy rises, the cognitive discipline required to maintain respectful treatment in group-based patrol is enormous.” (Applbaum 1996: 155) But this applies equally, although in different ways, to the ways members of the profiled group will perceive themselves, and the way the rest of society will perceive them.

Finally, *inequality* effects occur for several reasons, first and foremost of which is that social deprivation is certain to follow increased criminal supervision, both directly in the shape of incarceration, and indirectly in the shape of lessened job prospects and the lower income that this causes, damaged family structures and individual brutalization. But also as a result of the way both profiled persons and non-profiled persons may respond rationally by seeking segregation. If my chances of being subjected to a stop-search increase when I give my colleague a lift, because she possesses trait T, then it makes sense for me to give a lift instead to a different colleague who, like myself, does not. Inequality is thus likely to increase as profiling concentrates burdens that would otherwise have been more widely distributed on a particular group. Concentration is also likely to create added burdens, such as those that follow from rational strategies of segregation, and from cumulative effects of deprivation, since the groups targeted are likely to be groups that are already socio-economically deprived given how strongly socio-economic deprivation normally correlates with the target trait of criminality.

Overall I share Annabelle Lever’s concern that: “...we should expect racial profiling to exacerbate racism in society at large, even in apparently unrelated areas such as housing, transport, employment, and entertainment. [...] It is likely that racial profiling discourages black people from living, travelling and working in white neighborhoods, especially at night, and so compounds residential and occupational segregation. It is likely that it discourages black people from joining the police, and so perpetuates a damaging public perception of the police as hostile to black people. It is likely that racial profiling obscures the fact that most violence is intra-racial,

rather than inter-racial and committed by a minority of people, whatever their color.” (Lever 2007: 23; cf. also Loury 2005; Harris 2002; Harcourt 2004)

Two factors which will heavily influence these costs are the publicity of the profiling and the social character of the group profiled. The publicity of profiling is likely to have a strong impact on deterrence, alienation and stigmatization, whereas the social status of the group is most likely to influence alienation, stigmatization and inequality promotion. Generally speaking, publicity, that is, the recognition that profiling is being applied, will increase the redistribution of deterrence, aggravate alienation and deepen stigmatization. On the other hand, if hypothetically a form of profiling could somehow be implemented without anybody realizing that it was in use, it would be unlikely to incur these costs.

Simultaneously, the social character of the group profiled will play an important role. Most of the concerns regarding profiling focus on its use to target groups that are both highly socially salient, centrally racial, ethnic and religious groups where “...perceived membership of it is important to the structure of social interactions across a wide range of social contexts” (Lippert-Rasmussen 2007: 386), and which are socio-economically deprived and/or marginalized. Obviously, targeting some groups will have a greater impact than targeting others, say, “blacks” versus “bankers” (for financial crimes), both because the identification of the group *as* a group is stronger, and because bankers are not commonly perceived nor do they perceive themselves to be in a vulnerable position or the victims of historic injustice at the hands of the state. Or, as David Wasserman aptly describes how most persons might experience it: “It is far less demeaning to incur suspicion because of one’s accidental or voluntary associations than because of one’s race. It is, for example, merely annoying to become an object of close police scrutiny because one attended an opera at which gunshots were fired. It is profoundly demeaning to be subject to close police scrutiny because one is a black male youth.” (Wasserman 1996: 117) These qualifications make it even more difficult to generalize about the quantity of costs attached to different types of profiling. Such issues must be settled on a case-by-case basis, taking the social context into consideration.

In conclusion, we should be aware that even in situations where the beliefs upon which profiling are based are true it is unlikely that the attitudes it will promote are rational. It is far more likely that applying a form of profiling, particularly one involving highly socially salient proxy-traits, no matter their accuracy, will also promote unfounded biases. Secondly, we should expect profiling to have unintended bad consequences, such as the segregation of residential areas

and the loss of deterrence, which are neither emotional nor personal, but reflect rational strategies for responding to the situation of profiling. Such costs must be included when weighing the pros and cons of a profiling practice.

To sum up, profiling involves the following benefits and costs:

Benefits and costs of profiling

Increased benefits through higher ratio of true positives	Anti-deterrent effects
	Alienation effects
Decreased costs through lower ratio of false positives	Stigmatization effects
	Inequality promoting effects

Furthermore, building on the definitions previously established, we may now define the conditions of “minimally justified profiling:”

Costs, benefits and minimal justification of profiling:

1. Police practice PP using proxy-set S has probability P_I of identifying a randomly selected person A with trait(s) T_1 (and $T_2, T_3 \dots$) as having committed a criminal act;
2. Police practice PP using proxy-set S_{+T} has probability P_{I+T} of identifying a randomly selected person A with trait(s) T (and $T_1, T_2, T_3 \dots$) as having committed a criminal act ;
3. True positives produce benefits B_{true} and costs C_{true} ;
4. False positives produce benefits B_{false} and costs C_{false} ;
5. Using S_{+T} instead of S adds costs (anti-deterrence, alienation, stigmatization, inequality) C_{+T} .

Profiling for trait T on police-practice PP is *minimally justified iff*: The marginal benefits of using S_{+T} outweigh the costs, i.e. $((P_{I+T} - P_I) \cdot B_{true} + ((P_{I+T} - P_I) \cdot C_{true}) + ((1 - P_{I+T}) - (1 - P_I)) \cdot B_{false} + ((1 - P_{I+T}) - (1 - P_I)) \cdot C_{false}) \geq C_{+T}$

While hardly innovative in terms of normative theory, nor indeed particularly refined in terms of probabilistic thinking, I believe that this captures the essentials in a concise format. And it explains our conflicting intuitions, as I attempted to illustrate them initially by the opposed pull of *Chicago Customs* and *London Metropolitan*. We want at once to focus the efforts of the police only on those who deserve it, because this minimizes the costs imposed on the innocent, and to not focus the efforts of the police in a way that incurs costs by alienating, stigmatizing and promoting inequality. Weighing these concerns is a central challenge in the moral assessment of police-work. Having thus set out the framework for assessing profiling in consequentialist terms, let me turn at last to the reasons why I think racial profiling is unlikely to be justifiable.

Three Challenges – Foundation, Valuation and Application

In the following I will sketch out three challenges that racial profiling faces. Together, I believe they constitute at least a prima facie case against racial profiling, as well as illustrations of why proponents will have a hard time meeting this prima facie challenge. The first concerns the relation of racial profiling to potentially unjustified police practices, as the glossing over of unsound foundations, the second concerns the uncertainties of the costs and benefits of racial profiling, but particularly the benefits, and the third concerns the difficulty of properly applying racial profiling and the diminished efficiency likely to arise from cognitive and psychological biases.

The Foundation Challenge

Given that consequentialism is committed to holding an action to be morally wrong unless it is at least as good as all available alternatives, i.e., the maximisation of good consequences, the requirement of being minimally justified may seem to fall trivially short of the goal. As Applbaum observes: “Among the set of search strategies with positive net benefits, some are better than others. There may be a more refined strategy that is more efficient... [...] If a more refined search strategy is available, not to use it is inefficient.”²¹ (Applbaum 1996: 146)

²¹ Cf. also e.g. Lippert-Rasmussen who suggests that “unalloyed racial profiling” relies, among other criteria, on the fact that: “...no alternative, and equally or more effective, way of doing this is feasible.” (Lippert-Rasmussen 2006: 191-192) and Howard McGary, who argues that targeting searches against members of one particular race, for the purpose of apprehending drug couriers, is not justified even if the correlations is non-spurious unless: “...there is a high probability that stopping only speeders from this group would maximize arrests for drug offenses”. (McGary 1996: 137)

Consider, however, under what circumstances it is true that there exists a potentially superior alternative. What are the given alternatives to profiling for T on PP? First and foremost, it is clear that profiling in the light of all relevant proxies is not only difficult, given the high uncertainty of many factors, but also requires resources in terms of time and effort, resources which will frequently not be available in the relevant situation, or which might be better spent in other ways in the relevant situation. A police-officer driving past a person on the street and considering whether to stop and search her does not have time to make the kind of complicated probabilistic or moral arithmetic outlined in this article. As Stephen Maitzen argues: "...a social policy involving a given level of [statistical discrimination] is justified if and only if the information-cost of further statistical refinement equals or exceeds the net social utility to be gained by such refinement." (Maitzen 1991: 26) Requiring police to use maximally efficient profiling must therefore be understood as the maximum effort compatible with optimizing the ratio of profiling benefits over selection costs. Secondly, an action need only be considered an alternative if doing it and profiling for T on PP are mutually exclusive. If they can be performed simultaneously and will both be beneficial then consequentialism will require us to do both, but even if one is clearly superior to the other this could not lead us to be morally required to perform the superior *rather than* the inferior action.

Keeping these considerations in mind, maximising profiling looks relatively simple because the only obvious alternative to profiling for T on PP is PP without profiling for T. It may be possible, though I am hard pressed to think of an example, that profiling for race, i.e., including race in the proxy-set which is applied to select targets for a practice, could preclude other forms of profiling, but it does not seem likely. Thus "a more refined search strategy" than racial profiling, if racial profiling is independently beneficial, will be racial profiling plus whatever other proxies refine the search strategy, rather than such other proxies without racial profiling. In this case it seems that the relevant alternative to racial profiling is not to profile for race, which means that minimally justified profiling *is* maximising.

This does not, however, get police practices off the hook. The obvious alternative to a police practice is also no police practice, i.e., the same situation without police subjecting anyone to the given practice. This is the alternative covered by the requirement of minimal justification. But the stricter requirement of consequentialism forces us to also explore whether there are other and more efficient ways of promoting the values which ultimately justify police practices, such as urban

renewal, improved educational opportunities, etc. Again, unless doing one prevents us doing the other, these do not count as alternatives in the pertinent sense, and this may not be true of many superior social policies aimed at the same purposes. But any police practice will also have a near-infinite number of variations where details in the practice are altered, which really are incompatible with the practice: we cannot ask officers to apply two different practices to the same situation simultaneously. Thus, when assessing practices we really do need to think carefully about maximising, not just minimal justification. Is this not a problem, however, for assessing practices rather than racial profiling? Perhaps not exclusively. For racial profiling, when applied to a practice that does not meet the requirement of maximisation, and in some cases not even minimal justification, risks supporting its continuation and prolonging a practice that is morally wrong.

As I attempted to show in *Chicago Customs* initially we can easily imagine practices where the costs imposed seem out of proportion to the benefits potentially produced. The answer in such situations could in one sense seem to be more profiling, rather than less, since a practice that imposes heavy costs can only hope to achieve minimal justification by having a high probability of apprehending offenders. But in realistic scenarios we might worry that something else will happen instead. The fact that racial profiling tends to target marginalized minorities may make it easier to ignore the fact that a practice is not, in fact, justified. In the case on which *Chicago Customs* is based this seems to have been one of the effects. It was clear to those subjected to the practice that there was no way it could be justified, but at least partly because they belonged predominantly to the twice marginalized group of black women their complaints fell on deaf ears for far too long. It took the lucky break of one subject encountering an interested reporter, and the resulting massive media attention, for authorities to take seriously the plight it had imposed on innocent travellers and reassess the practice. (Harris 2002)

Because racial profiling tends to redistribute burdens so that the costs of a practice are imposed on those least capable of protesting or refusing them, it risks supporting practices that – profiling or no – are not minimally justified. And because such police practices represent genuine social evils this risk must be counted against it. Polemically put, if racial profiling is built on the morally corrupt foundations of unjustified practices, and protects these from the renovation they urgently need, then we would be better off not profiling.

The Valuation Challenge

An important issue that I have thus far neglected to address concerns how exactly to weigh the various costs and benefits. Clearly, we are dealing with costs and benefits that are not immediately commensurable. Even supposing that we had accurate figures for the costs and benefits involved, how ought we to weigh a certain degree of stigmatization of a racial group and the expenditure of so-and-so much money and time against the prevention of a specific number of crimes? This is clearly a case of apples and oranges. Nor am I suggesting that there is – or at least that I know of – any easily applicable method of translating such disparate goods and evils into a common currency. First, because I have deliberately attempted to shape my analysis in a way that did not rely on any one particular value-theory I do not have resort to the intrinsic value(s) that these instrumental values serve, and which would be needed to assess commensurability. Second, and more profoundly, because even if we were to adopt a well-developed value-theory, e.g., hedonism with QALYs as the measuring unit, I do not want to claim that there is a simple and uncontroversial way of exchanging the different costs and benefits. Most of the work here will, I suspect, have to be done by intuitions, despite the serious methodological concerns we ought to have about these. (cf. Singer 2005) However, this is a problem which affects costs and benefits equally, and in a wider context a challenge for any ethical theory which grants consequences any weight. As such it does not constitute an argument for or against racial profiling in itself, but one difficulty we face when we want to assess it.

A different, but related problem is that the empirical facts of the matter are not easy to sort out. Obtaining reliable statistics about racial differences in the propensity for crime has turned out to be extraordinarily difficult, as has obtaining reliable statistics for the deterrent effect of police practices or criminal justice in general. (cf. Von Hirsch et al. 1999; Harcourt 2004) But the difficulties involved here, and the resulting uncertainties about the size of the instrumental benefits and costs, do not affect benefits and costs equally. On the benefits side, incapacitation is both measurable and predictable. We can tell exactly the incapacitory effect of a given practice and anticipate with relative confidence the effects of increased (or decreased) apprehensions resulting from adding or removing a form of profiling. Similarly, we can make relatively good assessments of any efficiency benefits in terms of resources and inconvenience. Both are easy to observe and the changes resulting from increased or decreased efficiency can be measured and anticipated. But not so for arguably the most important benefit: deterrence. This is both uncertain and difficult to

measure. Even assuming that racial profiling increases the efficiency of a practice, there is no guarantee that this marginal increase will lead to greater deterrence or easy way of finding out how much of a difference it makes if it does.

On the costs side of the equation effects are more certain. The marginal increase in punishment is as easy to observe and anticipate as the use of resources and incapacitation effects. Further, there can be little doubt that conditions will, in those cases where the correlation between race and crime justifies racial profiling, also make race a social phenomenon subject to the costs I have described. Some members of a profiled racial minority will almost inevitably feel ill-treated and respond with hostility towards the legal system, just as both some members of the minority and the majority will almost inevitably internalize negative stereotypes about the minority, or reinforce existing prejudices. And since profiling redistributes costs many of the inequality effects will occur simply as unintentional or rational responses to changes in the incentive structure. All these costs may be difficult to measure precisely, but there is little doubt that they will occur. Only anti-deterrent effects are subject to the same severe uncertainty that afflicts deterrence, and for the same reasons. Harcourt, however, summarizes his review of the data in that: "Given the paucity of the evidence on both relative elasticities and offending, the conclusion is tentative, but under these assumptions [relatively low general elasticity, slightly lower elasticity for minorities and slightly higher natural minority offending rates], racial profiling probably increases the profiled crime." (Harcourt 2004: 1371-1372)

Uncertainties, it seems, do not afflict both sides of the equation equally. Racial profiling is a form of profiling that carries heavy and plausible costs. Other forms of profiling will suffer from similar uncertainties on the benefits-side, but not need to weigh these against such heavy burdens of justification on the costs-side. With the benefits so uncertain, it may be hard to achieve minimal justification, a problem which is aggravated by the difficulty of maintaining the efficiency of racial profiling in its application.

The Application Challenge

The third and perhaps most serious challenge for racial profiling is that there is good reason to be sceptical about the possibilities of properly applying it. The challenge here can be traced to two factors: the difficulty of obtaining accurate and precise data about the correlation between race and crime mentioned above, and the further difficulty of using this data in an efficient

way because of cognitive biases, which is the topic here. Suppose for an instant that reasonably accurate and precise data could be obtained, what is the likelihood that it could be effectively applied? Not, I would suggest, as good as we would prefer. This is so because, as studies in decision theory and behavioural economics has demonstrated in increasingly convincing detail, we humans are not very good at coping with probabilities.

Of the many problems afflicting the application of probabilities in drawing inferences, one of the most serious is undoubtedly the failure to appreciate the influence of the base rate. Consider the following example, lightly modified from how it is presented in (Schauer 2003: 94-96):

Schauerville. A violent crime has been committed in a city where 85% of the population is white and 15% of the population is black. Crime-rates for the two races are known to be equal, but an eye-witness to this particular crime, e.g. the victim, claims that the assailant was black, a piece of evidence the reliability of which is estimated as 80%. Is it likely that the offender is black?

Our intuitions might lead us to combine the high probability of the witness being correct with the fact that since blacks make up such a small proportion of the population, and the witness identified the assailant as a black person, then it really must be so. But the answer is a counter-intuitive no. There is, in fact, a roughly 0.59 probability that the assailant was white. The assailant is white on average in 85 out of 100 cases and the observation only partially redistributes this base probability, because a much higher number of white offenders will be misidentified as black than vice versa.²²

²² By Bayesian inference $P(\text{white}) = 0.2 \cdot 0.85 / (0.2 \cdot 0.85 + 0.8 \cdot 0.15) = 0.5862$ Note also the way that unreliability will lead to exaggerated perceptions of the correlation of minorities with a trait. Suppose Schauerville is representative of average witness reliability in reporting offender race, i.e. the city has an 85% white and 15% black population, there is no correlation between race and criminality and witnesses are, on average, 80% reliable when reporting the race of the offender. Assume further, and probably unrealistically, that the reliability is the same no matter the race of the offender, i.e. that white offenders are as likely to be misidentified as are black offenders. The resulting *observed* crime rates work out at a shocking 29% crimes committed by blacks, and 71% committed by whites, which deceptively indicates that far from being equally criminal Blacks are roughly twice as criminal as whites. And this warping of the picture emerges, in this example at least, without any form of racial prejudice, simply because many more whites will be misidentified given that there are more of them. Put differently, there is a “regression towards the mean”-type effect in any situation in which populations are uneven and trait attribution unreliable.

In addition to the problem of insensitivity to the base rate at least two further distortions may importantly taint probability estimates. Since race is such a highly salient trait our cognitive biases make it highly likely that its predictive power will both be over-estimated, and that it will serve as the starting point of prediction after which an “anchoring” effect will take place, reducing the predictive power, and hence adjustment of the estimated probability, of other traits. Correcting for these biases is difficult, because the calculations are complex and the results counter-intuitive. Suppose that we expand the example to further illustrate the difficulties police will face when they try to use this evidence in practice:

Schauerville II. Police are now looking for the suspect, and considering whether to stop and search person A, a practice for which they estimate the probability of that person’s being an offender necessary to obtain minimal justification (i.e. $P_I \geq N$) as 0.02. Their proxy-set allows them to consider the following traits:

1. A is black; blacks are 15% of the population; evidence gives a 41% chance that the offender is black;
2. A is male; males are 50% of the population; 90% of violent crimes are committed by men;
3. A is young, youth are 20% of the population; 50% of violent crimes are committed by youth;
4. A looks nervous; 5% of persons observed by police look nervous; 50% of offenders observed by police look nervous;
5. A is in the vicinity of the crime-scene; 100.000 people are in this part of the city; there is one offender among them.

Ought police subject A to the practice? Obviously, even if they were lucky enough to somehow have such exact figures, police officers cannot be expected to work out mathematically the probability of A being the offender, but will have to rely on their intuitive estimate of how likely the traits combine to make it. Given the low requirements – police can afford to search 49 innocents for every offender – and the combination of relatively many suspicion-raising factors it may seem that the answer is yes. In fact, A’s probability of being the offender is only about 0.0012, well below the

0.02 threshold where subjecting him to the search could be expected to do more good than harm.²³ Importantly, none of the problems I've sketched here represent flaws on behalf of individual agents, in the sense of psychological distortions such as racial biases and the like. As Daniel Kahneman and Amos Tversky put it: "These biases are not attributable to motivational effects such as wishful thinking..." but rather represent commonly used cognitive heuristics that while "highly economical and usually effective" also lead to "systematic and predictable errors". (Kahneman and Tversky 1974: 1130,1131)

We must not ignore the influence of psychological distortions, though. Added to the systemic weaknesses of our thinking about probabilities described above is the spectre of racism which, much as we would prefer the situation to have improved, is likely to at least occasionally haunt even modern police departments. The problem may not be as bad everywhere as e.g. suggested by the OSJI-report, which holds that the opinion expressed by a Spanish police officer that "all murders are related to immigrants (as are) 90 percent of drug crimes and gender violence" reflect commonplace perceptions in Europe. (Quoted in: Open Society Justice Initiative 2009) But the question is still how widespread such prejudices are, not whether they exist. The existence of police prejudices does not in and of itself affect the justifiability of profiling. Like police abuse it constitutes a problem independently of profiling and the obvious solution is to work at eradicating such prejudices, whether or not profiling is being applied. But as long as it exists institutional racism is worrying also because it deepens our concerns over the efficiency of profiling. Racist police officers are likely to further over-estimate the importance of racial differences in crime-rates, weakening the beneficial effects of using even a non-spurious correlation between race and crime.

Consider how systematic over-estimation of the effect of race as a proxy-trait would affect profiling. As illustrated in Profiling 101a the marginal benefits of profiling is the difference it makes in the average probability of a practice apprehending offenders. It does so by helping police estimate the individual probabilities of persons and target only those persons where $P_I \geq N$ and, in situations where they must choose, the person with the greatest probability of being an offender. Overestimating the effect of race will lead to police subjecting some persons of the profiled race

²³ Start by assigning A the 1:100.000 probability of being the offender based on the number of people in the vicinity of the crime scene. We can then adjust this by Bayes' theorem for being black $((0.00001 \cdot 0.41)/0.15)$: 0.0000273, being male $((0.0000273 \cdot 0.9)/0.5)$: 0.0000492, being young $((0.0000492 \cdot 0.5)/0.2)$: 0.000123 and looking nervous $((0.000123 \cdot 0.5)/0.05)$: 0.00123.

whose probability is actually lower than N to the practice, and it will lead them to sometimes prefer persons of the profiled race to persons of a different race when the latter would accurately have been estimated to have a higher probability of being offenders. In both cases this will decrease the resulting frequency of true positives and increase the frequency of false positives.

How bad the results of overestimation are for profiling depends on two factors. The first is the degree of overestimation. The second is the predictive power of race as a proxy-trait for crime; the higher this is the greater the margin of error. If the predictive power is sufficiently strong and misestimation is sufficiently small then profiling will maintain its benefits. But as the first grows weaker or the second greater, the benefits of profiling will diminish. Eventually, if the two factors cross a certain threshold, profiling will *reduce* the benefits of the practice in addition to still carrying its full costs.

Although the degree of inefficiency inherent to the application of a given proxy will, of course, depend on the specific situation in which the proxy is applied, it seems to me that there is good reason to be sceptical that proxies that, like race, are at once costly to use and have relatively limited predictive power will sustain enough of their marginal benefit to meet the minimal justification.

Conclusion – the Art of the Unseen

Throughout the course of this article I have sketched both what I take to be the consequentialist framework for assessing racial profiling, by separating profiling from the practices in which it is applied, reviewing the costs and benefits of each and establishing the conditions of minimally justified profiling. Doing so enables us to tackle the best arguments for racial profiling head on, as these tend to rely exactly on the proposed benefits of targeting racial groups when and where membership of such a group correlates with crime. However, I have argued that the situation is considerably more complex than is sometimes assumed, and in particular that racial profiling is a form of profiling which comes with a steep price attached. Profiling for race will, in realistic scenarios where race correlates with crime, be considerably more costly than most other profiles, because race is invested with such importance as a marker of identity. Alienation, stigmatization and inequality can all be expected to accompany the use of racial profiling. Given this background, I have sketched three challenges, which I think combine to at once make the case for racial profiling difficult to argue, and together establish a *prima facie* case that racial profiling will not, in many or

perhaps most circumstances, achieve minimal justification. As always when weighing consequences the balance ultimately depends on empirical facts, facts which in the case of racial profiling are extremely difficult to ascertain with any degree of certainty. In lieu of such ultimate clarification we will have to resort to intuitions and rough assessments of the strengths and weaknesses of the policies we put into practice. My suggestion is only that racial profiling, it seems, does not look promising.

If we accept this conclusion and decide to abandon racial profiling until such time as its justifiability can be better decided, the next challenge remains, for what does it mean for a police officer to not profile for race, and how would we determine whether she had or not? Setting prejudice aside, once police know that racial differences in propensity to crime exists it may be hard if not impossible to ignore them when estimating the probability that a person is an offender. This problem is compounded by the fact that such estimations are informal or intuitive, for how will an officer know whether or not her estimate has subconsciously included race as a proxy-trait? She might try to compare her estimate with a hypothetical estimate of the same person if that person had been a different race, but such thought experiments are bound to be tricky, the conclusions vague and subject to all sorts of biases themselves – the reluctance to admit to oneself that race plays a subconscious part in one’s reasoning not the least. These are very human problems, not problems restricted to instances of racism, and we should expect to find them not just in situations where officers are subject to racial prejudice, but in all police-work in communities with racially diverse populations. The real challenge, I would suggest, is therefore practical or psychological, rather than moral. The question remains how to ensure that racial profiling does not take place; how, that is, to help officers practice the art of the police, and see not that which it is useless that they should see.

Acknowledgments

I have presented drafts and parts of this paper at the 2009 meeting of the Nordic Network for Political Theory, Oxford University’s James Martin Advanced Research Seminar, Copenhagen University’s Centre for the Study of Equality and Multiculturalism (CESEM) “Discrimination in the Criminal Justice System” seminar and the 2010 Racial Profiling conference in Copenhagen. I am grateful for many useful comments and suggestions that I have received on these and other occasions from among others Paul Bou-Habib, Michael Boylan, Angelo Corlett, Claus Festersen, Claus Hansen, Bernard Harcourt, Jakob Holtermann, Nils Holtug, Klaus Frovin Jørgensen, Xavier Landes, Annabelle Lever, Kasper Lippert-Rasmussen, Morten Ebbe Juul Nielsen, Jeffrey Reiman,

Simon Rippon, Mathias Risse, Kira Vrist Rønn, Anders Sandberg, Julian Savulescu, Inge Schiermacher, Walter Sinnott-Armstrong, Stine Hansteen Solhøy, Pablo Stafforini, David Wasserman and Niklas Olsson Yaouzis. I owe particular thanks to Jesper Ryberg and Roger Crisp for very detailed and helpful comments.

References

- Appelbaum, Arthur Isak. 1996. Racial Generalization, Police Discretion, and Bayesian Contractualism. In *Handled with Discretion*, ed. John Kleinig, 145-157. Lanham: Rowman & Littlefield Publishers.
- Cox, Chana Berniker. 1993. On Michael Levin's "Responses to Race Differences in Crime". *Journal of Social Philosophy* 24: 155-162.
- Gross, Samuel R., and Debra Livingston. 2002. Racial profiling under Attack. *Columbia Law Review* 102 (5): 1413-1438.
- Harcourt, Bernard. 2004. Rethinking Racial Profiling: A Critique of the Economics, Civil Liberties, And Constitutional Literature, and of Criminal Profiling More Generally. *University of Chicago Law Review* 71 (4): 1275-1381.
- Harris, David A. 2002. *Profiles in Injustice*. New York: The New Press.
- Politiken. 2009. Politiet finder ingen våben ved visitationer. 24 August.
- Kahneman, Daniel, and Amos Tversky. 1974. Judgment under Uncertainty: Heuristics and Biases. *Science* 185: 1124-1131.
- Kennedy, Randall. 1997. *Race, Crime and the Law*. New York: Vintage Books.
- Lever, Annabelle. 2007. What's Wrong with Racial Profiling? Another Look at the Problem. *Criminal Justice Ethics* 26 (1): 20-28.
- Politiken. 2010. 18.000 visitationer gav 300 våben.
- Lippert-Rasmussen, Kasper. 2006. Racial Profiling versus Community. *Journal of Applied Philosophy* 23 (2): 191-205.
- Lippert-Rasmussen, Kasper. 2007. Nothing Personal: On Statistical Discrimination. *Journal of Political Philosophy* 15: 385-403.
- London Metropolitan Police Authority. 2010. Stops and Searches Monitoring Mechanism May 2010. London: London Metropolitan Police Authority.
- Loury, Glenn C. 2005. Racial Stigma: Toward a New Paradigm for Discrimination Theory. In *Understanding Poverty*, eds. A. Banerjee, R. Benabou, and D. Mookherjee, 654-665. Oxford: Oxford University Press.
- Maitzen, Stephen. 1991. The Ethics of Statistical Discrimination. *Social Theory and Practice* 17: 23-45.
- McGary, Howard. 1996. Police Discretion and Discrimination. In *Handled with Discretion*, ed. John Kleinig, 131-144. Lanham: Rowman & Littlefield Publishers.
- Open Society Justice Initiative. 2009. Ethnic Profiling in the European Union: Pervasive, Ineffective and Discriminatory. New York: Open Society Justice Initiative.
- Phillips, Coretta , and Ben Bowling. 2002. Racism, Ethnicity, Crime and Criminal Justice. In *The Oxford Handbook of Criminology*, eds. Mike Maguire, Rod Morgan, and Robert Reiner, 579-619. Cornwall: Oxford University Press.

- Risse, Mathias, and Richard Zeckhauser. 2004. Racial Profiling. *Philosophy & Public Affairs* 32 (2): 131-170.
- Schauer, Frederick. 2003. *Profiles, Probabilities and Stereotypes*. Cambridge: The Belknap Press of Harvard University Press.
- Singer, Peter. 2005. Ethics and Intuitions. *The Journal of Ethics* 9: 331-352.
- Statistics on Race and the Criminal Justice System 2008/09. 2010. ed. Ministry of Justice.
- Thomas, Laurence. 1992. Statistical Badness. *Journal of Social Philosophy* 23: 30-41.
- Von Hirsch, Andrew, Anthony E. Bottoms, Elizabeth Burney, and P-O. Wikström. 1999. *Criminal Deterrence and Sentence Severity*. Cambridge: Hart Publishing.
- Wasserman, David. 1996. Racial Generalizations and Police Discretion. In *Handled with Discretion*, ed. John Kleinig, 115-130. Lanham: Rowman & Littlefield Publishers.