

Racial Profiling: A Reply to Two Critics

Mathias Risse

John F. Kennedy School of Government, Harvard University

September 24, 2006

1. Risse and Zeckhauser (2004) seek to delineate the shape of the moral debate about profiling, to provide conceptual clarification and to think through some arguments.¹ Since then, Lever (2005) and Lippert-Rasmussen (2006) have offered thoughtful responses to our article, questioning some of its central points. While both of them offer insights that advance this debate (some of which we will adopt), this article aims to defend the main arguments in Risse and Zeckhauser (2004).

Since “racial profiling” is a pejorative term that was introduced to criticize police practices, one needs to be careful stating what one is assessing morally. Three issues are often confused. The first is the use of race as an information-carrier for investigative purposes; the second is police abuse; and the third is the “disproportionate” use of race in profiling. Often, cases that motivate discussion of profiling involve obvious instances of police abuse, or of overuse of race, so as to render superfluous any reflection on conditions under which race might be used legitimately in police tactics. It may be true that, given current conditions, it is impossible to use race as an information carrier without encountering massive problems of police abuse. Yet even if so, an investigation into conditions under which race could in principle be used is worth our while, to understand what are good and bad reasons for or against doing so. Our concern, then, is to assess the use of race as an information carrier in police investigations. We define

¹ The present article was prepared for a session on racial profiling at the 2006 Eastern APA in Washington, D.C. It will be thin on references, but for those the reader should consult Risse and Zeckhauser (2004).

racial profiling as “any police-initiated action that relies on the race, ethnicity, or national origin and not merely on the behavior of an individual.” We ask: Are such actions justified under conditions that could plausibly arise? (It is in this sense that we ask about the “legitimacy of profiling.”) Unlike other definitions of profiling,² ours leaves it open whether there are such circumstances.

Profiling is peculiar in two ways that make an assessment of its legitimacy different from such assessments in other contexts, especially employment discrimination. First, we are concerned with a fundamental public good (security), and second, situations where profiling will be used are those in which investigators must make quick decisions about (say) whom to search, or where large numbers of people are involved, situations, that is, in which not much information is accessible other than an individual’s appearance. In other cases (e.g., hiring or admissions) more information is available. These two points need to be kept in mind in what follows. The reason why profiling poses a serious moral problem is precisely because of these points.

There is a range of different situations for which questions about the legitimacy of the use of race arise, and whose particular features require equally particular considerations. To mention three paradigmatic cases: there are scenarios where police take measures employing race and ethnicity that seek to apprehend individuals who have committed specific crimes; there is racial, ethnic, or nationality screening at airports,

² In a typical approach, Ramirez et al. (2000) define profiling as “any police-initiated action that relies on the race, ethnicity, or national origin, rather than the behavior of an individual or information that leads the police to a particular individual who has been identified as being, or having been, engaged in criminal activity.” This definition *contrasts* (a) the use of race, ethnicity, or national origin with (b) the use of an individual’s behavior or information that helps apprehend somebody who has been identified as having been engaged in criminal activity. This suggests that profiling serves purposes *other than* apprehending criminals, imparting an aura of illegitimacy to it by definition. Moreover, writing that contrast into the definition suggests that *either* one uses race, ethnicity, etc., *or* one uses specific information on suspicious activity, namely information about an individual’s behavior or information that leads to an individual.

widely discussed after 9/11; and there is investigations on highways that rely in part on racial criteria, with the goal of intercepting drug traffic, and investigations on streets with the goal of finding illegal weapons. These scenarios differ in the extent to which a crime or threat is immediate, the extent to which intrusions can be expected, and the magnitude of the imposition. Profiling is the more controversial (and the burden of justification the higher) the less immediate the threat that prompts it, the less one can expect to be subject to such measures, and the greater the burden imposed.

As a final preliminary, we assume, for the sake of the argument, that there is a significant correlation between membership in certain racial groups and the tendency to commit certain crimes. Moreover, we assume that police can curb crime if they stop, search, or investigate members of such groups differentially, and that such measures eliminate more crime than do other measures for equivalent disruption and expenditures of resources. We do not hold a view on these issues. Instead, these assumptions set the stage, and if either of them fails, our question no longer arises.

2. Lippert-Rasmussen seeks to identify a consideration he thinks our argument omits and that allows to identify the source of “widespread reservations” (p 203) about it. But first let me offer his excellent summary of our argument. The reader might benefit to hear our case summarized by a critic. He points out that we (“RZ”)

have tentatively argued that ‘in a range of plausible cases (...) racial profiling [is] morally justified. This is so, according to RZ, whether we presuppose a utilitarian or a deontological moral theory. In the case of utilitarianism, RZ point to the fact that crime prevention increases the quality of life for almost everyone. They also suggest that ‘the primary contributor’ to the harm involved in racial profiling – i.e., the chief cause of ‘resentment, hurt, and loss of trust among’ groups subjected to increased police attention as a result of such profiling – ‘is likely to be underlying racism or underlying socio-economic disadvantages, rather than

profiling as such.’ While these harms may be great, the incremental harm involved in racial profiling is comparatively small. In the case of deontological theories, the extra burdens imposed on some people by racial profiling are neither unfair, nor rights-violating. They are not rights-violating, because ‘profiling can be justified on legitimate grounds’, i.e., the need to curb crime; and in any case, unalloyed profiling³ as such ‘does not contribute to [an] oppressive relationship involved in pejorative discrimination against one racial group by another. The extra burdens are not unfair, because whether ‘functioning reciprocity’ obtains or not, racial profiling is not unfair to those subjected to it. A society characterized by functioning reciprocity is one in which: ‘Differential burdens are imposed if five conditions are met: (1) it is required by the nature of the public good. (2) The imposition occurs through a suitable democratic process that gives proper weight to the interests of all involved. (3) The social importance of the good is widely acknowledged. (4) State interference with individual lives entailed by the imposition of differential burdens is within the limits of what citizens can reasonably be expected to bear – in particular the imposition of the good does not involve the violation of widely acknowledged rights of individuals. (5) It is widely acknowledged that, similarly, differential burdens are imposed for other public goods as well, and that this differential imposition of burdens for different goods works, by and large, to society’s advantage, although it may not work to every individual’s advantage.’ Where functioning reciprocity obtains, to deny that racial profiling is a fair means of achieving the collective good of reduced crime on the grounds that it imposes differential burdens on those subjected to it would, in effect, be to imply that collective goods can hardly ever be brought about in a way that is fair. But this is implausible. Even if functioning reciprocity does not obtain, racial profiling is fair to those subjected to it, because even they (or at least almost all of them: those who are law-abiding and perhaps even some of those who are not) will be net-beneficiaries. At any rate, this is likely to be the case when most crimes committed by members of a particular racial group are against other members of the same group. (p 192/93)

The only qualification I would like to add is that, towards the end, we are talking about conditions under which profiling is justified, instead of claiming that these conditions are generally satisfied.

³ Unalloyed profiling is profiling of the following sort: “First, the police activities are based in the right way on justified, true beliefs about crime rates in different racial groups. Second, the groups identify profiling [sic] enables the police to catch more perpetrators of crime, prevent more crimes and thus, reduce the overall incidence of crimes, i.e., it promotes the aims of the police relating to crime. Third, no alternative, and equally or more effective way of doing this is feasible. Fourth, the profiling is appropriately deployed, i.e., applied in all and only in those situations where this will lead to better crime related results. Finally, useful information concerning the crime rates in groups other than racial groups is treated no differently than equally useful information concerning the crime rates in different racial groups” (p 191f).

Lippert-Rasmussen focuses on the idea, “crucial” to our argument (p 193), that, overall, profiling is beneficial to those subject to it. Otherwise, he says, the utilitarian argument would be significantly weakened (there would be fewer situations where the balance of costs and benefits tilts in favor of using race), and the fairness argument would fail absent functioning reciprocity. He grants that profiling is beneficial to those subject to it (relative to a state that is otherwise unchanged but where there is no profiling), and explores what this fact contributes to its justification. The consideration missing from our argument, according to him, is that it is the whites who have created a situation in which profiling may well benefit the blacks. What is required is a change in the behavior of whites, rather than the imposition of extra-burdens on blacks. To spell this out, Lippert-Rasmussen assumes, for the sake of the argument

(i) that African-Americans are more likely to commit certain crimes than European-Americans solely as a result of the deprivation resulting from discrimination and unjust, racial inequality. Hence, if discrimination and unjust racial inequality were eliminated, the crime rates of European-Americans and African-Americans would converge over time. Suppose next, (ii) that all European-Americans could choose to act so that, in the long run at least, African-Americans would no longer suffer unjustly from discrimination and racial inequality. Suppose (iii) that given the existing discrimination and racial inequality, racial profiling will benefit African-Americans as well as European-Americans. It will benefit African-Americans because, although African-Americans will have to bear the costs of racial profiling, they will also enjoy the lion’s share of the benefits in the form of reduced crime, since African-Americans are more likely than European-Americans to be victims of crime. Suppose, finally, (iv) that relative to a state in which there is neither discrimination nor racial inequality, European-Americans generally benefit from discrimination and racial inequality in their favor, while African-Americans are generally harmed (a supposition that is consistent with the idea that in some respects European-Americans as well as African-Americans may benefit from the cessation of discrimination and racial inequality). (p 194)

Lippert-Rasmussen’s hunch is that people who agree with these assumptions (which he finds broadly plausible) will agree with his criticism of Risse and Zeckhauser (2004), and

it is in this way that he wants to identify those “widespread reservations” about our argument, whereas those who do not agree with them will reject his criticism.

Suppose whites point out to blacks that profiling is to their benefit. This presumably presupposes that they are actually concerned to benefit blacks. Yet says Lippert Rasmussen,

were European-Americans to act differently, discrimination against African-Americans and racial inequality would cease to exist. As a result, racial profiling would no longer be a net benefit to African-Americans since the crime rates of European-Americans and African-Americans would then converge, ensuring that racial proofing is ineffective in reducing crime. Hence, according to the normative premise of the argument [that police measures should be taken that work to benefit of African-Americans] – which presumably European-Americans accept, since otherwise they would not put forward the argument --- European-Americans should act differently. (p 196)

Call a “comprehensive” justification of a policy one such that, if that justification makes a factual assumption that a certain subgroup will act in a certain way, that way of acting is itself justified.⁴ Using this term, we can state Lippert-Rasmussen’s claim by saying that profiling cannot be comprehensively justified. This leaves a possibility for non-comprehensive justification. As will become clear, Risse and Zeckhauser (2004) are concerned only with such non-comprehensive justification, and Lippert-Rasmussen’s main contribution, I take it, is to point out that it is because of this that our argument leaves an important dimension of the debate unaddressed and thus encounters widespread reservations. While he also argues that our non-comprehensive justification fails, I will

⁴ Lippert-Rasmussen takes this notion from Cohen (1992). His argument is inspired by an incoherence objection Cohen raises against Rawls: If the rich take the standpoint of the difference principle, they say two things that do not make sense together: first, they say that a principle should be followed according to which things should be arranged to the benefit of the least advantaged. Second, they are asking to be given special incentives. But this second point runs contrary to the spirit of the first – after all, the least well off would benefit even more if the talented would not insist on incentives. Lippert-Rasmussen applies this to the profiling scenario: The claims here are that the concern is to benefit African-Americans; but then, what one should do is not to insist on profiling, but to change conditions so that blacks commit fewer crimes.

not discuss that claim. My discussion of his views on comprehensive justification, adding to Risse and Zeckhauser (2004), offers resources to rebut his objections in the non-comprehensive case, but I leave this for the reader to decide. Anyway, the non-comprehensive case is not where the greatest interest of Lippert-Rasmussen's article lies.⁵

3. Lippert-Rasmussen is right that one's attitude towards these assumptions will shape one's view of our arguments. To the extent, and in the form, that these assumptions are plausible, Risse and Zeckhauser (2004) accommodate them, and to the extent, or in the form, that they go beyond what we accommodate, they are implausible. Also, the notion of comprehensive justifiability must be revised to be of interest in this context. So while indeed profiling cannot be justified comprehensively given conditions (i) and (ii), this does not matter because both conditions and this notion of justifiability should be revised.

While quibbles with (ii) matter more for my argument, let me record two quibbles with condition (i) as well. First, while it must be true in some sense that, in the US, "whites" have discriminated against "blacks," what this means specifically will take on different senses for different individuals across generations. This is obvious, but it matters because Lippert-Rasmussen uses (i) to trigger the idea that it is the whites' fault that blacks display higher crime rates and that this diminishes their claim to the protection to which profiling might contribute. But given this differential involvement of whites in racist practices it is implausible that no, or few, whites now have a claim to the kind of protection towards which the use of race in police tactics is supposed to contribute.

⁵ It is actually inaccurate that the goal of profiling is to benefit blacks, which is what would be required for there to be an incoherence problem as Lippert-Rasmussen constructs it. The goal of profiling is to increase general security, and the question is how to distribute the burdens that must be incurred towards that end. The problem is better understood as a fairness issue rather than a coherence issue.

Second, what drives the inference in condition (i) to “if discrimination and unjust racial inequality were eliminated, the crime rates of European-Americans and African-Americans would converge,” is the word “solely” in the statement that “African-Americans are more likely to commit certain crimes than European-Americans solely as a result of the deprivation resulting from discrimination and unjust, racial inequality.” What is easy to agree to is condition (i) with “to a large extent” replacing “solely.” But this is too weak for Lippert-Rasmussen’s purposes because it does not contribute as much to the idea that whites do not have a claim to the kind of security to which profiling is supposed to contribute. However, “solely” introduces a mono-causality that will be hard to support. What motivates this strong version must be the idea that some racial, ethnic or national groups display higher crime rates than others only if, and then because, they have been disadvantaged in the past. Without wanting to engage in arm-chair social science, I think in this generality that is implausible.

To illustrate, recall Levitt and Dubner’s (2005) thesis that the reduction of crime widely observed in the US in the last 15-20 years is causally connected to *Roe v. Wade*, the 1973 constitutional ruling that legalized abortion. Roughly speaking, if abortion is legal, a lot of unwanted children remain unborn, many of whom otherwise would have been brought up under conditions parents had judged unsuitable for raising children. Such children have a higher likelihood of becoming criminals than others. So the fact that they remained unborn in the wake of *Roe v. Wade* is a possible cause of crime reduction. Suppose this is true. Suppose also that some racial, ethnic or national groups differ in their attitudes towards abortion. Suppose moreover that, despite these differences, these groups do not differ in terms of the way in which they would treat unwanted children,

and would want, or not want, children for the same reasons.⁶ If Levitt and Dubner are right, the group with a more conservative attitude towards abortion would have raised more criminals. There is no reason to suppose such differences in attitude towards abortion can only be explained in terms of past oppression. Such differences may be moral or cultural in rather innocuous ways. This is sketchy, but the point is that the term “solely” in condition (i) presupposes that nothing remotely like this account could hold. While that is conceivable, and while I have no way of arguing either way, one should not find condition (i) *obviously* plausible – it is not.

More serious, for our purposes, are my reservations about condition (ii), which states “that all European-Americans could choose to act so that, in the long run at least, African-Americans would no longer suffer unjustly from discrimination and racial inequality.” Note first that the envisaged changes are even more complex than the “mere” cessation of racist actions suggests. Bringing about a society not scarred by racism is not merely a matter of doing things differently *from now on*; it is also a matter of changing the current distribution of assets because racism can only stop shaping our society if such rectification has been performed. Moreover, those who have so far had less than they should have and have been treated differently from how they should have will not easily behave like people who have not been cheated in these ways even if rectification occurs. There are patterns of behavior that can be explained by ways in which people have been socialized in a racist society that will be hard to change even if circumstances do.

These points illustrate the magnitude of the task involved, and Lippert-Rasmussen could grant them. Yet the problem about condition (ii) is to make sense of the idea that

⁶ This condition is meant to make sure that the conviction that rules out or allows abortion is assumed not to translate into differential attitudes towards unwanted children.

we “could” change our behavior so as to bring about a non-racist society. It is only in a logical or physical and perhaps in a biological sense that we *could* all change in such a way as to create such a society. What matters is that there is no politically feasible sense in which we now “could” collectively so change our behavior. There is no action-guiding sense of “could” that makes this claim true, nor therefore one that would lead to blame if we failed. There are no blueprints for changing our ways that are also politically feasible in the sense that we could now embark on a well-understood democratic process leading us from where we are now to a society free of racism by implementing these proposals. We do not know how to do that within the confines of what is politically doable.

I do not mean that a society, once racist, is doomed to being racist forever. Still, transforming a racist society is highly non-trivial. Certain policies help more than others, but developing a non-racist society will be a matter of trial and error and of change over time that can only to some extent be shaped by policy choice. To push the point, lest this sounds unduly gloomy: If we could end racism in a sense of “could” connected to feasibility, we “could” do a lot: we could arrange traffic in such a way that nobody gets injured or killed; we could end all wars; generally, we could ensure people get through life with less pain and unhappiness, to the extent that this is determined by how we act in concert. More to the point, we could end crime and would not have to worry about profiling at all. In none of these cases is it impossible that societies can engineer such changes; yet unfortunately, nor is there a sense in which we “could” do it such that there is a reasonably clear connection to what we ought to do here and now such that if we do not act in that way, we can be blamed.

In *Civilization and Its Discontents*, Freud warns that, pace Marx and other social thinkers, resolving the conflicts shaping human living arrangements is not a matter of finding the right social order; instead, such conflicts arise from our instinctual nature. Even if one would not want to put it this way, it seems fair to say that we do not have the faintest idea of how to live together without crime. And at any rate, those agents typically concerned with assessing whether racial profiling should be used are agents whose actions might have a bearing on local crime statistics, but are very unlikely to have a bearing on society's overall racism. So condition (ii) is true, but only if the crucial term "could" is understood in a manner that is neither action-guiding nor allows for the assignment of blame. It is true only in a sense that does not help us in pondering what we had better do instead of profiling.⁷

4. Let us turn to Lippert-Rasmussen's notion of comprehensive justification. Recall that a policy is comprehensively justified only if it does not assume as a factual matter that a group will act in a way that is unjustified. This entails that crime prevention measures, no matter what their character, cannot be so justified. They all are based on the assumption that some people act in ways that are unjustified (which is why prevention is needed). Similarly, measures imposing damages on individuals because they have not taken precautions against unjustifiable behavior of others could not be so justified either. Some

⁷ Lippert-Rasmussen is right that Risse and Zeckhauser (2004) grant versions of (i) and (ii). We suggest that measures be taken to decrease discrimination "that could presumably be implemented with some success if a majority of the American population were to support them" (p 195); but suggesting such initiatives is different from assuming that people could just change their behavior in a way that society is transformed into a non-racist society. Lippert-Rasmussen argues that such changes should be undertaken instead of profiling; my point above is that we do not understand well enough what would need to be done to that effect, and have no feasible way of doing it, so as to dispense with profiling.

policies only come under consideration *as responses* to unjustifiable behavior, so their justifiability cannot depend on whether such behavior exists.⁸

The idea of comprehensive justifiability fails in such contexts because the idea of measures being justifiable in this manner seems to be that they can be part of a coherent set of policies applicable in an ideal world. In such a world crime prevention is not required because there would be no crime. But then the fact that profiling cannot be comprehensively justified ceases to be alarming. Being justifiable in that way is not a property that justifications for just any measure need to satisfy to be acceptable in the context in which they are discussed. Some measures only emerge as options in contexts where one cannot help but presume that some people will act badly.⁹

But while the idea of comprehensive justifiability does not deliver interesting results about profiling, there is a related idea that may, and that also may capture Lippert-Rasmussen's intuition better than the unnecessarily general idea of comprehensive justifiability. That intuition, I take it, is that policy P, which imposes special burdens on group G1, is not justified if it is because of the unacceptable behavior of group G2 that G1 is in the position in which it makes sense to impose special burdens on G1 to realize the goals pursued by P. Without worrying about how to spell this out in detail, let us call this sense of justifiability "conclusive justifiability," to have a different name.

⁸ It is important to see why the example is crime prevention and the responsibility for precautions, rather than crime itself. Measures against crime could be justified in conditional terms: "If people act in such and such a way, this and that should be done." No factual assumption would have to be made. But as soon as crime prevention must be justified, the assumption is that crime occurs.

⁹ Cohen, from whom Lippert-Rasmussen takes the notion of comprehensive justifiability, is aware of this. Anticipating the point made above, he says after introducing the term: "It follows, harmlessly, that penal policies adopted to reduce the incidence of crime lack comprehensive justification. The very fact that such a policy is justified shows that all is not well with society" (p 279, n 13). Right he is.

Should profiling be abandoned if it cannot be justified conclusively? This depends on the feasibility of policies that would improve the situation of G1 or the behavior of G2. This takes us back to the discussion of condition (ii). It is not enough that G2, logically, physically, or biologically, “could” change their behavior to render special costs for G1 redundant. The question is whether there is a feasible way of bringing this about. For any measure that cannot be conclusively justified there is a loss to imposing it, and hence a pro-tanto reason not to do so. Nevertheless, one may have to choose such a measure (e.g., profiling) anyway, because there is no feasible alternative given the conditions. One still might need to do as much as one feasibly can to change this situation, both by improving the situation of G1 and by altering the behavior of G2. But the general needs at stake (in our case, basic security) may be too important, and other measures (aimed at such improvement and alternation) too uncertain to take effect sufficiently fast for measures such as profiling to be expendable in the short run.

Risse and Zeckhauser (2004) argue that the view that profiling is justifiable under certain circumstances is consistent with measures improving the status of blacks. Lippert-Rasmussen is right that this claim should be strengthened. We should say that “to the extent that it is because of unacceptable differential access to privileges that blacks commit more crimes than whites, society has an obligation to take measures to improve their lot, and such measures involve both measures on their behalf and measures that change the behavior of whites.”¹⁰ However, what one should resist are these stronger

¹⁰ One might say that my response to Lippert-Rasmussen implies that too many collective failures to promote the good will be pardoned and condoned. To mention the obvious example: Could not the Germans have stopped Hitler, and are they not to blame for having failed to do so? But my main point here – and this is also reflected in the statement above to which this is a footnote -- is not that racism should simply be left to grow or disappear because there is nothing much we could and thus nothing much we should do about it; it is instead that racism cannot be changed easily, rapidly, and straightforwardly. That is, there is no sense in which we “could” change it here and now such that we are not entitled to resorting to

views: that profiling should be abandoned (a) because society is not entitled to imposing special burdens on blacks since it is the fault of whites that blacks are in a situation where that might advance the goal of fighting crime; or (b) because what society should do instead is just stop being racist; or (c) because profiling cannot be justified comprehensively or conclusively. Such views appear plausible only if one accepts conditions (i) and (ii) or endorses ideas about what it takes to justify such measures that do not apply to contexts where we need to inquiry about profiling.¹¹

measures such as profiling (which, again, serve the promotion of rather elementary public goods) if need be because we are not doing certain things to end racism. I think it is because this is my position that I need not introduce a detailed theory of what one can hold groups responsible for.

¹¹ From here, Lippert-Rasmussen moves on to the context of non-comprehensive justification. In that context, I would just like to comment on one point. Lippert-Rasmussen doubts that our utilitarian argument succeeds: “I suspect that from a utilitarian point of view, while a society free of discrimination and racial profiling is better than one with discrimination and racial profiling, the latter may well be better, for the reasons given by RZ, than a society in which there is discrimination but no racial profiling. This, of course, is significant if correct. However, from this it would not follow that racial profiling is non-comprehensively justified from a utilitarian point of view. That conclusion would follow only if (a) there were no alternatives to racial profiling that were both superior to such profiling and left the rest of society, including present racial inequality, unchanged; or (b) although there were alternatives of this kind, none was better than the option in which racial profiling is implanted along with other measures. RZ offer no reason to think that (a) and (b) are true and hence give us no reason to believe that racial profiling is non-comprehensively justified from a utilitarian point of view” (p 199). But recall the two assumptions that guide our inquiry: First, that there is a significant correlation between membership in certain racial groups and the tendency to commit certain crimes; and second, that police can curb crime if they stop, search, or investigate members of such groups differentially, and that such measures eliminate more crime than do other measures for equivalent disruption and expenditures of resources. Indeed, we do not show that profiling has the features Lippert-Rasmussen mentions; it is an empirical question whether it does, one we are not competent to settle, but the moral question that concerns Risse and Zeckhauser (2004) does not arise otherwise. Lippert-Rasmussen explores an alternative answer to his challenge: “It might be replied that they simply, and plausibly, assume that no genuine alternative involves a significant reduction in existing racial discrimination and inequality. But, first, this suggestion sits uneasily with RZ’s own list of measures to change that situation – e.g., Head Start, affirmative action programs, and reparations. To demand that it be shown that there are no logically possible alternatives to racial profiling that are at least as good would be to set the standards for utilitarian argument for racial profiling unreasonably high, but restricting oneself to a simple comparison of the status quo with and without racial profiling is to fail to set the standards high enough. Second, it would be odd anyway to assume that no significant reduction of existing racial discrimination and inequality is possible. For, surely, given the will, we could reduce racial discrimination and inequality.” (p 199) But at this stage, the response to this point is straightforward. Questions about profiling arise in specific crime-fighting contexts; and while the need to fight crime might well justify the use of race in police tactics, there are also reasons to change society in such a way as to make it true that race is not significantly correlated with a propensity to commit certain crimes (at any rate, to the extent that such a propensity is due to independently unjustifiable conditions, such as past oppression). Racial profiling should be thought of as a short- or medium-run measure, whereas such

5. Whereas Lippert-Rasmussen seeks to identify a consideration our account neglects, Lever engages one of our arguments directly and in depth, namely the utilitarian argument.¹² The crucial issue for an assessment of profiling from a utilitarian standpoint is how to think about the costs and benefits attached to it. Kennedy (1999) argues that the cost-benefit assessment cannot simply be done in terms of crime reduction (benefits) and inconvenience of those screened (costs); instead, the feeling of resentment among minorities, a sense of hurt, and a loss of trust in the police must also factor in. Once these costs are counted, Kennedy claims, an argument *against* profiling emerges.

Risse and Zeckhauser (2004) agree that these costs must be considered, but submit that profiling has such effects only against the background of a society minorities *already* perceive as racist. While profiling causes inconvenience and other harm, sometimes considerable, the primary contributor to resentment, hurt, and loss of trust is underlying racism or underlying socio-economic disadvantage, not profiling as such. If so, utilitarian considerations must factor in exclusively the *incremental* harm inflicted by profiling per se. This will be comparatively small if most of the overall level of harm that seems to be caused by profiling is plausibly ascribed to underlying causes. So in a range of plausible cases, utilitarian considerations will support profiling.

More specifically, we are offering two theses: the *expressive harm thesis*, and the *incremental harm thesis*. The harm attached to a practice or an event is “expressive” if it

measures to change society as a whole should be seen as long-term measures. Once one takes such a standpoint, such measures do not sit uneasily with profiling.

¹² Our article is not committed to this utilitarian stance. We also consider non-consequentialist arguments, but those generally do play much of a role in this response article because they are not central to either Lippert-Rasmussen or Lever.

occurs primarily because of harm attached to *other* practices or events. According to the expressive harm thesis, the harm attached to profiling is expressive: it occurs because of underlying racism. The expressive harm thesis concerns the way in which harm associated with profiling is caused, not its *subjective* importance or *subjective* extent; nor does it speak in any other way to the nature of the harm experienced. If the expressive harm thesis is true, the incremental harm thesis follows: the amount of harm associated with profiling should not count in the harm done through underlying racism and socio-economic disadvantage, but only what is actually added to that by the use of race in police tactics (i.e., the incremental harm associated with profiling). It is implicit in that thesis that this increment is much smaller than the harm Kennedy would associate with the use of race in police tactics.

To license the inference from the expressive harm thesis to the incremental harm thesis – the former of which, after all, is a causal claim, whereas the latter is a claim about how to do utilitarian calculations -- we need to add a principle saying something like this: A utilitarian policy maker should count as costs of policy P only costs specifically incurred through imposing P, not costs that also arise in the process of implementing P, but that are plausibly seen as caused by underlying socio-economic conditions, rather than by P itself. This principle seems plausible. So the burden of proof, I take it, lies on arguing for the expressive harm thesis. At the same time, the expressive harm thesis is of practical relevance because it entails the incremental harm thesis.

Contrary to what Lever suggests, our account does not require us to say those subject to profiling are “confused” (p 103) about their resistance to profiling. This impression should arise no more than it should for Kennedy (1999). Kennedy’s and our

concern is to offer a view on how to think about harm that should be counted as costs of profiling, and what these views share is a rejection of a naïve stance on that question that would think of those costs in terms of the inconvenience, or more generally, the decline in well-being caused on the spot by actual acts of profiling. There is more to an assessment of the costs of profiling than such an on-the-spot check. When Kennedy claims that “more” needs to be counted among the costs, he does not claim that those subject to profiling are confused. Nor do we make that claim when we follow his move, but suggest that what has been added must be evaluated differently. Our question arises from the third-person perspective of a utilitarian policy-maker who must decide how to assess costs and benefits. This does not involve any questioning of subjective experiences of those subject to profiling. It is, one may say, an exercise in moral accounting.

I will defend the expressive harm thesis against Lever’s objections. However, Lever also and independently criticizes the incremental harm thesis. One way of understanding her argument is to insist that even if we were right about the expressive harm thesis, we would underestimate the actual amount of harm caused by acts of profiling. To show this, Lever offers vivid descriptions of what it is like to be profiled. For instance, she reminds us that

[p]olice in the United States carry guns, and are known to use them. By the side of the motorway no one can really tell what is going on. A wrong move, the inability to hear or understand what is being said, a fit of coughing or a panic attack can all lead to violence and tragedy. Police have been known to mistake a black man gasping for air, or suffering from a heart attack or epilepsy, for someone trying to resist arrest or to attack them, and have then responded with what turned out to be deadly force. In those circumstances, I would surmise, fear would be an appropriate reaction to being stopped by the police, however innocent one is. [p 103f]

Lever is right that we pay insufficient attention to spelling out the nature of the incremental harm caused by profiling, and she might be right that, therefore, we sometimes make it sound as if this incremental harm were smaller than it is. On this point we stand corrected. But the crucial issue for us is to establish the expressive harm thesis as a way of thinking about the harm associated with profiling in the first place. The envisaged opponent is somebody who defends the sort of view proposed in Kennedy (1999), according to which much more must be counted as harm associated with profiling as such (and thus as costs in a utilitarian cost/benefit assessment of profiling) than we submit should be counted (given, again, that we agree with Kennedy's move that the costs of profiling cannot be assessed through an on-the-spot check of the well-being of those subject to the use of race in police investigations); and that can be true even if one accepts much of Lever's insistence that there is more to the incremental harm of profiling than we have allowed. So our main point holds up, or so I shall argue.

6. The expressive harm thesis, then, is a causal thesis, and as such an empirical claim whose proper verification is a question for the social sciences, provided the relevant data and conceptual tools are available. No such investigation is available, as far as I know, so for the time being we are limited to more speculative ways of thinking about the expressive harm thesis, that is, thought experiments and parallel scenarios.¹³

Lever's overall worry is that the expressive harm thesis downplays the importance of the kind of harm associated with profiling, and thus presents the use of race in police

¹³ This is not to say that this whole exercise here is merely unprofessional social science. One philosophical problem that comes along with the utilitarian argument is to assess just what should be considered the relevant costs and benefits in the first place. Generally, it lies in the nature of utilitarianism that an intricate mix of empirical and conceptual considerations enters into evaluations of moral questions.

tactics in an unduly favorable light. This worry is developed in different ways, and in great detail, so that the best way of addressing Lever's objection is to engage at least most of her points one by one, also in some detail. Lever's doubts begin by her questioning whether there really is a clear distinction between expressive and other harms to begin with. Most instances of harm, she thinks, will come with expressive dimensions; conversely, apparently clear instances of expressive harm will also plainly hurt. As Lever puts it, referring to our example of torture as a kind of harm that is not expressive:

Torture is painful no matter the reasons why it is inflicted. Still it does not follow that what we might call the neutral aspect of the pain of torture is the most significant thing about it. Hence I am skeptical that we can use Risse and Zeckhauser's idea of expressive harm to characterize and evaluate the harms of racial profiling, because the idea of an expressive harm seems, itself, to depend on our ability to make a variety of complicated causal and interpretive judgments. (p 100)

Lever comes to this assessment by discussing an Argentine scenario where the fact that certain political opponents were Jewish mattered for the decision to torture these, rather than other, people. Lever takes this to entail that Anti-Semitism must be accounted for in any assessment of the kind of harm inflicted by such torture, which she thinks is a matter of difficult empirical, rather than a priori, judgment.

Recall that we define expressive harm as harm occurring *primarily* because of the presence of other practices. This qualification is important because it is plausible that the nature of any kind of pain or other harm is *to some extent* determined through the presence of other practices. Just about any action is embedded into a social context, and thus how actions benefit or harm others cannot be seen entirely independently of what else has happened to or been done by the persons involved. Nevertheless, there is a point to distinguishing forms of harm in which what matters *most* is *not* this kind of social

embeddedness, and thus in which the presence of other practices is not what needs to be accounted for primarily, and forms of harm for which this is the case. Snubbing is a paradigmatic example of the latter, and torture is a paradigmatic example of the former. In many other cases that are less clear it will be a matter of, as Lever puts it, “complicated causal and interpretive judgments” to determine the extent to which the relevant harm is due to social embeddedness of certain actions.

Lever’s Argentine torture scenario leaves me unconvinced. I do not know what it is like being subject to torture in a context where ethnic discrimination matters for the decision whom to torture although torture is actually not inflicted for ethnic reasons; but I doubt that we should have any difficulty in saying that “the neutral aspect of the pain of torture is the most significant thing about it” (p 100), where that neutral aspect is infliction of pain, rather than Anti-Semitism. As long as we can make sense of the idea of elementary human needs, we will be able to make sense of a notion of harm that is not primarily due to the social embeddedness of certain actions, and therefore also of a notion of harm that indeed is.

Or consider Lever’s example of the “flasher:”

“Flashers,” or men who display their penises to others, especially unsuspecting women and girls, are generally thought of as harmless, and their victims are encouraged to laugh off the behavior, and to see any shock, horror or fear on their part as quite disproportionate to the harm involved, and as itself a bit of a joke. The point, of course, is that no physical injury occurred: this is not rape; it is not even a beating or a mugging. (...) Insofar as it is felt to be more traumatic, this might be, to use Risse and Zeckhauser’s framework, because it is an expression of living in a sexist society, in which women are raped and murdered (...) Seeing the reaction this way would be nicer than seeing it simply as irrational behavior (...) But just because these stories seem nicer and less patronizing from the outside does not make them so. After all, we could put the matter another way: the harm caused by flashing, in a sexist society, is not an expression of some other harm, but one of the forms that sexism takes. It is often experienced as frightening and threatening, out of proportion to the harm caused by unexpectedly seeing a penis,

because it *is* frightening and threatening, and *meant to be so*. (...) Hence I think we should be wary of the claim that the harms of racial profiling are expressive of racism, rather than instances of racist harm. (p 99f)

This case is a good occasion to remind the reader of a distinction internal to the notion of expressive harm that Risse and Zeckhauser (2004) draw. Harm may be expressive if an event or practice is a *reminder* of other painful events or practices, or if one event or practice becomes a *focal point* for events or practices, a symbol of structural disadvantage or maltreatment. In addition to whatever harm the practice itself causes (and of course we concede that there is such harm), the focal point becomes associated with harm attached to such disadvantage, and that harm plausibly accounts for the lion's share of the harm associated with that practice.

Lever's flasher-case is a good example of an instance of expressive harm being a focal point. There is nothing about this notion that rules out that the relevant act (in this case, the unwanted display of a penis) is frightening and meant to be frightening. The work this notion is supposed to do is to explain *why* that act is perceived as frightening *to such an extent*. As soon as she says that flashing "is often experienced as frightening and threatening, out of proportion to the harm caused by unexpectedly seeing a penis, because it *is* frightening and threatening, and *meant to be so*," Lever implicitly acknowledges the need for a notion such as expressive harm. The two thoughts she connects with "because" can precisely *not* be so connected unless it is also explained why the harm is experienced as frightening in a manner she concedes is "out of proportion" to what seems to be happening on the face of it. There is nothing about the notion of expressive harm that would commit us to advising women who have fallen prey to a flasher to "laugh off" this experience, or to think of them as irrational for being frightened, nor is this true of

profiling. What we claim is that the notion of expressive harm is useful, and probably necessary, to make sense of situations such as this flasher-scenario.

To add one more thought to this explication of the notion of expressive harm: In the flasher-quote Lever insists that the harm caused by flashing should be seen not as an expression of other harm, but as “one of the forms that sexism takes.” The worry seems to be that our account classifies the harm caused by profiling as an epiphenomenon, rather than as the harm that must be associated with one of a set of racist practices, and thereby makes the harm caused in this manner appear less important than it is. Risse and Zeckhauser (2004) address that point, insisting that there is nothing about the notion of expressive harm that implies phenomena classified as carrying that kind of harm are epiphenomena *rather than* some of those practices that constitute, say, racism (p 151f). Therefore, pointing out that profiling should be seen as one of the forms racism takes also does not entail that the harm associated with it cannot be expressive. Merely insisting on that point does not make it right, of course, but I hope this point gets some confirmation from the account of expressive harm above.

7. The distinction between expressive and other forms of harm, then, seems valid. But does it help in an assessment of racial profiling? As mentioned before, we offer two kinds of argument, thought experiments and parallel cases. Recall the thought experiments:

[I]magine the closest possible world to US society except that there is no racism. Race-caused disparities in economic or educational attainment do not exist, and practices such as race-related police abuse are unknown. In such a society, we conjecture, using race for investigative purposes would not be considered offensive and would not trigger resentment, hurt, or loss of trust in law enforcement. Consider also a possible world that differs from ours in that racial profiling has been effectively banned. Ours could become such a world following, say, a sweeping Supreme Court ruling outlawing racial profiling. As a result, the

levels of resentment, hurt, and loss of trust among minority group members, we conjecture, would *not* be significantly lowered. Simply stopping the practice of profiling would do little to change society's underlying racism and thus little to alter the *attitudes* that lead to police abuse and also promote various forms of racism in other segments of life. (p 146)

Lever (p 96) points out that no evidence is given for the Supreme-Court scenario. While it is true that we do not say more about it, this claim continues to strike me as plausible. It is not absurd to imagine the Court passing a verdict as envisaged here, and to image for this to happen although not much else has changed in our society, or is being changed along with such a judgment. This might be a matter of concerns about profiling making their way up to the Court, and of there being a particular constellation of judges who rule accordingly. Yet unless many other things are changed at the same time, it is not plausible that the socio-economic arrangements that form the background to continuing racism, as well as the attitudes of individuals, will change dramatically. When discussing Lippert-Rasmussen's view that we all "could" start acting differently and end racism, I pointed out the magnitude of change that would be involved here. How could this kind of Court judgment contribute all too much to this?

Consider, then, the parallel scenarios we envisage:

For instance, men between 15 and 40 commit a disproportionate share of violent crimes. Thus screening them (at the expense of controls on elderly women) is often justifiable. Hardly anybody (including those men) finds such measures offensive. No hurt is connected to membership in that group. Similarly, the white community did not object to the disproportionate attention given to whites – mistakenly in retrospect -- in the search for the Washington-area sniper in 2002. As another example, Ben Gurion Airport employs strict screening mechanisms for visitors exiting Israel. Security personnel decide in interviews whom to search. One criterion that tends to trigger a search is if the visitor spent time in Arab areas. Again, it seems that this measure is not offensive, given the security problems emerging from such areas. (This comparison is relevant only as long as we talk about tourists: if we are talking about Arabs, it becomes question-begging.) Consider another case that makes our point. More crimes are committed in summer. Thus it is reasonable to have more police patrols in July and August,

targeting people who are disproportionately out and about then. Examples of this sort abound. To be sure, these are not, strictly speaking, cases of profiling as defined in section 2. Still, they make the same point: treating people differently in accordance with some of their ostensible characteristics (in this case whether they are out disproportionately in July and August) often does no harm beyond the inconvenience. Harm is greater only if the characteristic for which a person is targeted is *independently associated* with (and thus makes such targeting a focal point of) the harm brought to the fore by the relevant action (say, a search). (p 148)

Lever is not convinced. She offers two points in response. First, she suggests that, in the Sniper case, white people might have thought they were screened as potential witnesses or as residents of suburban Washington. They might never have felt they were targeted in virtue of being white, and might in fact not have been targeted in virtue of being white. This is possible (though absent closer empirical investigation our version of it is possible as well), just as it will be plausible that, in situations in which indeed a population is targeted not on racial, but on residential grounds, blacks might think they are targeted on racial grounds. But be that as it may, the general point here seems plausible enough and does not stand and fall with our ability to offer conclusive empirical verification for one particular scenario. That point is that in some cases people do mind being treated in terms of a group-membership and in others they do not, and which is which depends on the social esteem in which that membership is being held.

Here is Lever's second response to these parallel cases. She says:

Assume, though, that Risse and Zeckhauser are right, and that white people do not object to the profiling of whites *because* no harm is attached to membership in the group "white." What does that prove? It gives us some reason to suppose that the racial profiling of black people is harmful because of the racism of a racist society. What it does not show, and cannot show, is that the harm of such profiling is largely expressive. (p 102f)

What is puzzling here is the transition from the second to the last to the very last sentence. Indeed, what Zeckhauser and I argue is that these parallel scenarios (alongside

the thought experiments) give us “some reason to suppose that the racial profiling of black people is harmful because of the racism of a racist society.” This is what we mean when claiming that “the harm of such profiling is largely expressive,” and we enlist these parallels and thought experiments to provide “some reason” in support of that view. Perhaps Lever disagrees with the implication we take this to have, namely, that, therefore, when we go through the utilitarian cost/benefit considerations, the costs of profiling must be discounted. That is, what is harmful about profiling because profiling also is a focal point of underlying racism cannot all be summed up under costs of “using race for police purposes.” Perhaps Lever’s point is to question the claim that the expressive harm thesis entails the incremental harm thesis, that is, to question the inference from the claim that “it is largely because of underlying racism that the use of race in police tactics is harmful” to the claim that “what should be counted as costs of the use of race in police tactics is merely the incremental harm done by that use of race, not the overall harm done by racism, which is largely responsible for negative consequences such as loss of trust, resentment, etc.”

One thing to say to make this inference more plausible is this. Suppose utilitarians went about assessing the costs of any practice that might be a reminder or focal point of underlying racism by counting as costs all consequences in whose causal explanation appear both that practice and underlying racism. Utilitarians will need some way of comparing quantities of harm, and I assume they have some such way. If they proceed as I sketched, the overall amount of harm done by racism, no matter how large, will be massively overstated. This is so because plenty of separate practices would be associated with underlying harm, but the overall harm caused by all these practices must be

aggregated from the harm associated with these practices individually, which leads to an overall much bigger amount of harm than what actually is associated with the underlying racism of society. In other words, there would be less harm associated with racism in society overall than with the aggregate of individual practices that constitute social interaction in society, which is absurd. It would take large efforts in the metaphysics of harm to make this train of thought precise. But the point is that utilitarians will need to be able to make sense of such reasoning in some form, and so if something like this is roughly correct, a denial of the incremental harm thesis given the expressive harm thesis would indeed lead to this sort of overestimation.

8. This is most of what I have to say in defense of the utilitarian argument in Risse and Zeckhauser (2004). I have argued that the distinction between expressive and non-expressive forms of harm is valid; that the two sets of arguments we offer in support of the expressive harm thesis hold up; and that the expressive harm thesis does entail the incremental harm thesis. These are the main plunks of our utilitarian argument. However, let me add some more responses to other important passages from Lever (2005).

Lever generally thinks ours is a “deeply implausible picture of the relationship between racism and racial profiling,” because it

requires us to imagine [1] a racist society in which racism has no role in explaining the choice of racial profiling over other ways of responding to racial disparities in crime that are, quite likely, a legacy of racism. Similarly, it asks us to suppose [2] that the profiling of black people can be carried out fairly and respectfully in a society that can still be characterized as racist. Only on these assumptions could objections to racial profiling be understood as objections to racism *rather than* to profiling; and only on these assumptions could the harms of racial profiling be identified with racism *instead of* racial profiling. These assumptions are far-fetched. The expressive harm thesis, therefore, depends on an

implausible picture of the relationship between racism and racial profiling. [p 96; makers “[1]” and “[2]” added]

But neither [1] nor [2] is entailed by anything Risse and Zeckhauser (2004) say. Of course, in a racist society decisions about which police tactics to use will often be racially motivated; and possibly, in such a society, it may not be feasible to make sure such measures are executed fairly. (But could not improved training take care of many of the problems?) As I pointed out before, possibly in practice profiling should not be used *for such reasons*; but then the reason for not using it is that the police cannot be trusted, not that there is something per se wrong in the use of race in police tactics.

As a final passage from Lever, consider the following, in which she argues that the expressive harm thesis

makes profiling seem more peripheral to racist habits of thought and action than it is, and draws attention to the way that profiling reflects racist attitudes, institutions and habits while obscuring its contribution to them all. Racial profiling publicly links black people with a tendency to crime. For that reason alone, it is likely to exacerbate the harms of racism. However scrupulous the police, racial profiling is likely to remind blacks, all too painfully, that odious claims about their innate immorality and criminality justified their subordination in the past, and still resurface from time to time in contemporary public debate. So too, racial profiling will likely perpetuate, as well as reflect, white tendencies to draw invidious and complacent racial distinctions, and exacerbate unmerited indifference and hostility to the legitimate interests of black people. (p 97)

More specifically, Lever identifies mechanisms through which profiling fosters racism:

Profiling increases the likelihood that whites will think of blacks as importing crime into their supposedly crime-free neighborhoods; it discourages black people from traveling and working in white neighborhoods, especially at night; and it provides a seemingly endless source of events, rumors and stereotypes for popular entertainment along ‘cops and robbers’ lines, happily obscuring the fact that the ownership of all sorts of horrendous weapons is perfectly legal in the United States. Thus, the racial profiling of black people has the unfortunate effect of publicly associating blacks with criminality, while downplaying any responsibility of white for that violence, or for the violence that they, themselves, commit. (p 97f)

At this stage, I need to remind the reader of the different paradigmatic cases of profiling distinguished above: profiling immediately connected to a crime; preventive screening of the sort that happens at airports; and screening on highways or city-streets in connection with drug-traffic or illegal weapons. In all these cases, there is some specific target pursued by the use of race. The kind of scenario Lever envisages seems to use race without any specified purpose whatsoever. We all agree that there are costs connected to such use; the disagreement is about how to think about these costs. At any rate, these costs need to be compared to sizable benefits, and on our way of thinking about costs and benefits it is very hard to see how any positive verdict would emerge for random screening of blacks passing into white neighborhoods *merely* because they do such passing. We envisage a scenario in which there is some specific occasion or background problem that renders the use of race a plausible device in crime fighting. It would then be a matter of police-community relations to communicate the reasons for such use accordingly, and to carry out the measures appropriately.

9. To put our inquiry into proper perspective, it is appropriate to close these responses to Lippert-Rasmussen and Lever with the same thought with which Risse and Zeckhauser (2004) concluded. Our argument that profiling might be justified requires that other people would have a complaint if we did not strive strongly to reduce crime. These harms are perceived on an uneven scale. Those hurt by profiling are readily identified; those protected by measures that fight crime are statistical figures. But it is indeed to the extent that such people are protected, and numerous, that arguments in support of profiling gain strength. People thus protected, many of them white and many not, do not

lose their claim to protection because of past or present racism; they do not lose their claim to protection because there is some, at any rate politically currently not feasible, sense in which society “could” transform into one in which the empirical assumptions of our inquiry do not hold (in particular that there is a correlation between membership in certain minorities and certain crimes); and they do not lose their claim because the use of race in police tactics cannot be justified comprehensively or conclusively.

These were the points argued against Lippert-Rasmussen. The points argued against Lever concern how, from the perspective of a utilitarian policy maker, one should compare the claims of those who demand protection with the claims of those who would have to carry some of the burdens that come with providing public safety. On that score, I have argued that the relevant cost-benefit calculation should be guided by the expressive harm thesis and the incremental harm thesis. I have argued that the distinction between expressive and non-expressive harm is valid; that the two sets of arguments Risse and Zeckhauser (2004) offer in support of the expressive harm thesis hold up; and that the expressive harm thesis does entail the incremental harm thesis.

Literature

Cohen, G.A. 1992. “Incentives, Inequality, and Community.” In: G. Petersen (ed.), *The Tanner Lectures on Human Values* 13. Salt Lake City: University of Utah Press

Kennedy, Randall. 1999. "Suspect Policy: Racial Profiling Usually Isn't Racist. It can help Stop Crime. And it Should be Abolished." *The New Republic*: September 13 & 20, 1999

Lever, Annabelle. 2005. "Why Racial Profiling is Hard to Justify: A Response to Risse and Zeckhauser." *Philosophy and Public Affairs* 33: pp 94-110

Levitt, Steven and Stephen Dubner. 2005. *Freakonomics. A Rogue Economist Explores the Hidden Side of Everything*. Willam Morrow

Lippert-Rasmussen, Kasper. 2006. "Racial Profiling Versus Community." *Journal of Applied Philosophy* 23 (2): 191-205

Ramirez, D., McDevitt, J, and A. Farrell. 2000. *A Recourse Guide on Racial Profiling Data Collection Systems: Promising Practices and Lessons Learned*. Report prepared by Northeastern University with support by the U.S. Department of Justice.

Risse, Mathias, and Richard Zeckhauser. 2004. "Racial Profiling." *Philosophy and Public Affairs* 32: 131-170