

Is Racial Profiling a Legitimate Strategy in the Fight against Violent Crime?

Neven Sesardić¹

Received: 21 August 2017 / Revised: 15 December 2017 / Accepted: 2 January 2018
© Springer Science+Business Media B.V., part of Springer Nature 2018

Abstract Racial profiling has come under intense public scrutiny especially since the rise of the Black Lives Matter movement. This article discusses two questions: (1) whether racial profiling is sometimes rational, and (2) whether it can be morally permissible. It is argued that under certain circumstances the affirmative answer to both questions is justified.

Keywords Racial profiling · Discrimination · Police racism · Black lives matter · Bayes's theorem · Base rate fallacy · Group differences

1 Introduction

The Black Lives Matter (BLM) movement is driven by the belief that the police systematically discriminate against blacks. If such discrimination really occurs, is it necessarily morally unjustified?

The question sounds like a provocation. For isn't it obviously wrong to treat people differently just because they differ with respect to an inconsequential attribute like race or skin color? Indeed, racial profiling does go against Martin Luther King's dream of a color-blind society where people "will not be judged by the color of their skin, but by the content of their character." The key question, however, is whether this ideal is achievable in the real world, as it is today, with some persistent statistical differences between the groups.

Figure 1 shows group differences in homicide rates over a 29-year period, according to the Department of Justice (DOJ 2011: 3):

As we see (the red rectangle), the black/white ratio of the frequency of homicide offenders in these two groups is 7.6 (i.e., 34.4/4.5). Such a large difference is a robust result that comes up in other studies as well. Actually, the real B/W ratio is even higher

✉ Neven Sesardić

¹ Zagreb, Croatia

	Percent of—			Rate per 100,000	
	Victims	Offenders	Population	Victims	Offenders
Total	100%	100%	100%	7.4	8.3
Age					
Under 14	4.8%	0.5%	20.0%	1.8	0.2
14-17	5.2	10.6	5.8	6.6	15.0
18-24	24.4	37.5	10.6	17.1	29.3
25-34	28.7	28.0	15.6	13.7	14.9
35-49	22.8	17.1	21.1	8.0	6.7
50-64	8.9	4.9	14.7	4.5	2.7
65 or older	5.1	1.6	12.3	3.1	1.1
Sex					
Male	76.8%	89.5%	48.9%	11.6	15.1
Female	23.2	10.5	51.1	3.4	1.7
Race					
White	50.3%	45.3%	82.9%	4.5	4.5
Black	47.4	52.5	12.6	27.8	34.4
Other*	2.3	2.2	4.4	3.8	4.1

Fig. 1 Group differences in homicide in the U.S. (1980–2008)

than 7.6 because the DOJ categorizes Hispanics as “white,” though Hispanics are significantly more likely to commit homicide than non-Hispanic whites.

Although these race differences in violent crime rates are large, real and stubborn, there is a lot of obfuscation, verbal manipulation and outright dishonesty in discussions of this topic.

A striking illustration of this “red herring” strategy can be found in a 2003 document issued by the DOJ (2003: 1):

[Racial profiling] is premised on the erroneous assumption that any particular individual of one race or ethnicity is more likely to engage in misconduct than any particular individual of another race or ethnicity.

So, federal law enforcement officers are instructed by the DOJ that racial profiling of blacks should be rejected because it relies on the “erroneous assumption” that any black person is more likely to engage in misconduct than any white person.

Well, has there been anyone in the contemporary world who was tempted to accept that “assumption”? Such a person would have to believe, for instance, that Barack Obama is more likely to engage in misconduct (or crime) than, say, the leader of Hells Angels. Virtually no functioning person could possibly believe that.

Racial profiling has never been based on the silly claim that *any black* is more prone to crime than *any white*, but rather on the richly and repeatedly corroborated hypothesis that blacks *as a group* commit violent crimes at a higher rate than whites. It is from this fact that racial profiling draws its justification.

In another DOJ document (2014: 3) we read:

[I]t is patently unacceptable and thus prohibited under this Guidance for law enforcement officers to act on the belief that possession of a listed characteristic signals a higher risk of criminality.

Notice that the DOJ doesn't say that the belief that some of the listed characteristics are connected with a higher risk of criminality is false. It cannot say this primarily because, at least with respect to one of the listed categories (gender), it is too glaringly obvious to everyone that being a man is linked with a higher risk of criminality than being a woman. But with race, again, the DOJ cannot deny group differences in the risk of criminality because it would contradict itself (see Fig. 1). So, the Guidance does not state, let alone establish, that the belief in question is false, only that law enforcement officers are prohibited *from acting on it*.

Why? Because this is "patently unacceptable". In another place the same document denounces profiling as irrational, claiming that it is not "smart and effective" law enforcement.

Basically, there are three main questions about racial profiling. First, is it rational? Second, is it morally permissible? And third, is it real and widespread? Critics, as well as many defenders, of the police agree that the answer to the first two questions is "No". There is a lot of disagreement about the third question.

I will focus mainly on the first two questions, and I will try to show that the answer to both of them is "Yes". With respect to the third question I will only briefly criticize attempts of some police supporters to demonstrate the nonexistence of police profiling.

Section 2 will address the first question, the rationality of profiling, while section 3 will deal with the second question, the moral permissibility of profiling. Section 4 will offer a brief conclusion.

2 Rational or Irrational?

In order to discuss the possibly different treatment of whites and blacks by the police it is best to start with some real data. Unfortunately, accessing this kind of information is made difficult by the widespread practice of many government sources of lumping together data for whites and Hispanics. For this reason I will look at crime statistics provided by two police departments—in [New York City](#) and [Chicago](#)—that keep a separate category for Hispanics and in this way make possible the crucial comparison between blacks and (non-Hispanic) whites.

Figure 2 (below) shows the arrest data for murder and non-negligent manslaughter among whites and blacks in the two cities in the last five years for which public records are available. For some reason the data for Chicago have not been publicized after 2010.

Since the number of blacks and whites in the population of Chicago is roughly the same, the B/W ratio in the table tells us directly how much more likely someone will be arrested for murder or non-negligent manslaughter if he is black than if he is white. In New York City, though, the number of whites is much higher than that of blacks. So in that case the percentage ratio should be corrected for the difference in size of the two groups in order to arrive at the true ratio of the arrest probabilities. And this explains why, in the above table, the B/W ratio for NYC is higher than what would be obtained by just dividing raw percentages in the column "Blacks" with those in the column "Whites".

Now which of the two stances is advisable for the police under the circumstances—color blind or color conscious? To see the difference between the two approaches let us

	Year	Blacks	Whites	Black/white ratio
Chicago (raw numbers)	2006	234	11	21.3
	2007	274	14	19.6
	2008	255	15	17.0
	2009	273	10	27.3
	2010	190	8	23.7
NYC (percentages)	2011	55,5	6,2	15.9 (corrected)
	2012	51,4	9,2	9.9 (corrected)
	2013	54	6,5	14.9 (corrected)
	2014	61,8	2,9	37.9 (corrected)
	2015	59,1	6,8	15.4 (corrected)

Fig. 2 Arrests for murder and non-negligent manslaughter in Chicago and NYC

introduce some additional assumptions which will help us derive the diverging implications of the two strategies. The all-important starting point will be that the arrest rate for murder and non-negligent manslaughter is around 15 times higher among blacks than among whites. (This is a conservative estimate, given that the arrest ratio is either around that level or higher in 9 out of 10 years in the whole 10-year period.) Let us also assume that the arrest rates roughly correspond to the rates those crimes are committed by the two groups. This seems to be a reasonable assumption, but I cannot provide detailed support for it here. The alternative hypothesis that the higher arrest rate of blacks is due to police discrimination faces many problems (e.g. see Beaver et al. 2013; Mac Donald 2016a: 151–162). Anyway, my goal is not to analyze empirical data but rather to illustrate the logic of the argument for discrimination and evaluate it.

Now we are ready to confront the central question of whether, under the above assumptions, the police have good reasons to treat a black person differently from a white person even when *specific* evidence for potential criminality is the *same* in both cases.

In order to facilitate the calculation of relevant probabilities we need to introduce the only part of the equation that is still missing, namely the information about specific evidence. For that purpose imagine there is a piece of evidence (E) which is much more often present among those who are guilty of murder or non-negligible homicide (M) than among others ($\sim M$).

To be specific, suppose that, *among both whites and blacks*, the probability of E, given M, i.e. $p(E|M)$, is 0.3, whereas the probability of E, given not M, i.e. $p(E|\sim M)$, is 0.0003. Now since E is much more frequent among Ms. than among $\sim Ms.$, the presence of E in a person will markedly increase the probability that the person has M.

See Figure 3 for a mathematical formula (called “Bayes’s theorem”) that gives us the probability that a given member of either group who possesses suspicious characteristic E is also guilty of murder or non-negligent manslaughter (M):

The expression $p(M)$ refers to what is called the *prior* probability of M, i.e. the probability that a person is M based only on the information about his group membership (in our case, black or white). Recall that in our scenario $p(M)$ is 15 times higher

among blacks than among whites. For ease of calculation let us assume that $p(M)$ among blacks is 0.15% and that, consequently, the same probability among whites in that hypothetical scenario will be 0.01%. (Nothing hinges on the specific numerical values assigned to these probabilities, which are introduced here just for the purpose of teasing out implications of the high B/W ratio in one specific scenario).

The probability we are interested in, $p(M|E)$, depends on several elements in the equation in Fig. 3, but all these elements are identical for blacks and whites, except $p(M)$ and $p(\sim M)$. The latter is of course immediately derivable from the former, because $p(\sim M) = 1 - p(M)$. This means that in our scenario *any* difference in $p(M|E)$ between blacks and whites will be due *only* to the difference between blacks and whites with respect to $p(M)$.

In other words, if the probability of a black suspect being a murderer turns out to be higher than for a white suspect (with the *same* specific evidence being present in both cases), this will solely be the result of the difference in prior probabilities of M in the two cases.

Let us arrive at the probabilities we want by presenting the situation in terms of frequencies, which is known to be much more intuitive and easier to follow than the approach that uses mathematical equations about probabilities.

First, consider 100,000 randomly chosen blacks from our scenario and see what proportions in this group are expected to belong to different groups defined by characteristics M and E . See Figure 4 for how it would look.

The probability $p(M|E)$ is obtained by dividing the frequency of $E \& M$ (the upper framed entry) by the frequency of all E , i.e. $E \& M + E \& \sim M$ (the sum of the both framed entries). So the probability that a randomly selected black person with characteristic E is also M is 0.6.

Figure 5 gives the same graph for whites.

The difference is striking. The probability of a person with suspicious characteristic E being a murderer is 0.6 if he is black, but 0.09 if he is white. And this difference is *exclusively* the result of different prior probabilities of M among blacks and whites.

Now imagine a policeman patrolling an area of Chicago or New York City and confronting an unknown person with characteristic E . Suppose, further, that the policeman is aware of the relevant frequencies of E and M among blacks and whites.

Will the policeman's behavior be different depending on whether the other person is black or white? It is hard to see how it could be the same. The level of caution, preparation for self-defense, readiness to draw a gun and shoot (if necessary) must be increased in the more threatening case. It would simply be irrational to behave in the same way in the two cases if in one case the probability of mortal danger is less than 0.1 and in the other case it is 0.6.

Putting it in general terms, prior probabilities must be taken into account in order to obtain the probabilities of the event in question. And it is this probability, based on *total evidence* (i.e. specific evidence *plus* prior probabilities), that will be an all-important factor determining which action is appropriate under the circumstances.

$$p(M|E) = \frac{p(M) \times p(E|M)}{p(M) \times p(E|M) + p(\sim M) \times p(E|\sim M)}$$

Fig. 3 Bayes's theorem




Fig. 4 Frequency of M among E, among 100,000 blacks

This is all just elementary probability and decision theory. Paying attention to prior probabilities (or “base rates”, as they are sometimes called) is not only permissible; it is actually *mandatory* for a rational person. Ignoring them is a logical fallacy known as the base rate fallacy (Tversky and Kahneman 1982).

But although it is mathematically demonstrable that neglecting base rates is a failure of reasoning, there are situations where there is tremendous pressure to commit that very fallacy. Psychologist Philip Tetlock talks about “*forbidden* base rates”:

Forbidden base rates refer to any statistical generalization that devoted Bayesians would not hesitate to enter into their probability calculations but that deeply offends a religious or political community. The primary obstacle to using the




Fig. 5 Frequency of M among E, among 100,000 whites

putatively relevant base rate is not cognitive, but moral. In a society committed to racial, ethnic, and gender egalitarianism, forbidden base rates include observations bearing on the disproportionately high crime rates and low educational test scores of certain categories of human beings. Putting the accuracy and interpretation of such generalizations to the side, people who use these base rates in judging individuals are *less likely to be applauded for their skills as good intuitive statisticians than they are to be condemned for their racial and gender insensitivity*. (Tetlock et al. 2000: 854—emphasis added)

People are notoriously afraid of being labeled as “racists”, which is probably the most lethal moral accusation today. Therefore, as Tetlock observes, those treading in this minefield try hard to cleanse themselves of “forbidden thoughts” by expressing “moral outrage” and “support for ostracizing and punishing deviant thinkers”, which culminates in “blanket refusal even to contemplate certain thoughts”. So, in a curious twist, what should otherwise be a praiseworthy effort to avoid a logical fallacy (“base rate neglect”) becomes a thought crime. Neglect base rates in this kind of case, or else!

To summarize:

- (1) Incidence of violent crime is much higher among blacks than among whites.
- (2) Different levels of violent crime in the two groups constitute different base rates.
- (3) Racial profiling is the act of taking these base rates into account in estimating posterior probabilities.
- (4) Not taking the base rates into account is irrational.
- (5) Taking the base rates into account is rational.

Therefore:

- (6) Racial profiling by the police in the fight against violent crime is rational.

Statement (1) is the only empirical premise and it is not seriously disputable. Information provided by the DOJ (Fig. 1), [data from the FBI](#), and from various police departments (e.g. Fig. 2) all corroborate (1). Statement (2) and (3) are true by definition, i.e. they are based on the meaning of the expressions “base rates” and “racial profiling”. Everything else follows from the elementary mathematics of probability.

Although racial profiling in this context can be a matter of life and death, it becomes such a taboo that its prohibition sometimes overcomes even the basic instinct of survival and self-preservation. A dramatic example was a female police officer who was brutally beaten and almost killed by a black man but did not use her gun because she feared accusations of racial profiling. A colleague who visited her at the hospital reported (*Chicago Tribune*, October 7, 2016): “She looked at me and said she thought she was gonna die, and she knew that she should shoot this guy, but she chose not to because she didn’t want her family or the department to go through the scrutiny the next day on national news.”

Although people may not dare ever mention publicly the forbidden information about different prior probabilities of violent crime between blacks and whites, they will tend to keep this in mind privately and react accordingly, as illustrated even in the case

of the black civil rights activist Jesse Jackson, who famously said, as reported in the *New York Times* (December 12, 1993):

There is nothing more painful to me at this stage in my life than to walk down the street and hear footsteps and start thinking about robbery—then look around and see somebody white and feel relieved.

Nevertheless, the huge majority of those who write about police shootings are adamant that racial profiling is irrational. The only disagreement between them is on whether they think the police do engage in racial profiling or not. The proponents of Black Lives Matter and many others believe that racial profiling by the police is endemic. Many of those who defend the police claim that the police are color blind, but they accept the premise that different treatment of whites and blacks would be both irrational and ethically unacceptable. (Recall the quoted DOJ statement that racial profiling is *not* “smart and effective”.)

One of the most outspoken police supporters, Heather Mac Donald (much of whose work I admire), often speaks of “the *myth* of racial profiling”, and “specious” and “alleged” racial profiling attributed to the police. Furthermore, after defining “soft profiling” as “us[ing] race as one factor among others in gauging criminal suspiciousness” (Mac Donald 2003: 10), she says (emphasis added):

[S]tate police commanders should eliminate any contribution that soft profiling may make to that perception [that blacks are stopped partly because of their race], unless *strong evidence* emerges (as it has not so far) that soft profiling has had an *extremely high success rate* in drug interdiction (16-17).

Notice how high the standard of acceptance is set for permitting soft racial profiling (which basically means paying attention to the base rates): according to Mac Donald, soft racial profiling should be forbidden unless there is *strong* evidence that it has an *extremely high success rate* in fighting crime. So, even a high success rate is *not* enough; it has to be *extremely* high. And with “extremely high” remaining undefined, will any success ever qualify?

In contrast, Mac Donald has a much lower standard of acceptance for the empirical claim, which she endorses, that the police do not engage in racial profiling against blacks. In a YouTube video (Mac Donald 2016b) which has been watched by more than 1.7 million viewers, she says that Harvard economist Roland Fryer concluded in his recent paper (Fryer Jr 2016) that in a study conducted in Houston “blacks were 24% *less likely* to be shot by officers than whites even though the suspects were *armed or violent*” (Mac Donald’s emphasis). There are several problems with her presentation of Fryer’s findings.

First, she fails to mention that Fryer’s study is only a working paper and that it was not peer reviewed (i.e. not validated by other experts). Second, her viewers are not informed that there have been other, *published* studies that contradict Fryer’s conclusion. Third, the data that Mac Donald presents with so much emphasis about the allegedly lower probability of blacks being shot by police are described by Fryer himself as “statistically *insignificant*”, which means there is no sufficient evidence to treat the putative difference as real, rather than a mere statistical fluke. And fourth, the

viewers are not told that as soon as Fryer's article was posted on the internet, and before Mac Donald's video was released, many experts and well-informed commentators immediately raised serious concerns about Fryer's methodology and expressed skepticism about drawing from his paper the kind of conclusion that she defends (Lind 2016; Simonsohn 2016; LaCapria 2016; Gelman 2016; Balko 2016). For all these reasons her short video is highly tendentious and misleading.

In addition, Mac Donald again does not report Fryer's findings correctly in her recent article:

In a data set comprising officer shootings from Dallas, Austin, Houston, Los Angeles and six Florida counties, [Fryer] found that officers were 47 percent less likely to discharge their weapon without first being attacked if the suspect was black than if the suspect was white. (Mac Donald 2016c)

In fact, Fryer does not say that officers *are* 47% less likely to shoot a black suspect, but only that "officers *report* that they are 47.4% less likely to discharge their firearms ... if the suspect is black" (emphasis added). A huge difference, no doubt. It is not just that what police officers say about their lack of anti-black bias *may* not be true. Furthermore, we know that, given the tremendous public outrage against racial profiling, police officers may have a good reason to make false claims that cast themselves in a good light. And even if the officers indeed entirely ignore prior probabilities, this would not mean that their behavior is reasonable, nor even that they regard it as reasonable. It might be that they engage in it just in order to avoid probable public condemnation.

We cannot end this section without considering an interesting argument against profiling. According to this argument, *even granted* that there are large race differences in base rates of violent crime, police would be more effective if they ignored the base rates in their decision making. The argument is based on a psychological, rather than a logical, premise.

The basic idea is very simple. Because of the salience of characteristics like race, people might ascribe much more importance to race than it actually has. For this reason, the argument goes, it might be necessary to prohibit any use of race, despite its indisputable informational value. As Frederick Schauer puts it:

[W]e might worry, *on efficiency grounds* as well as on moral ones, about the overuse of race and ethnicity and the consequent underuse of other relevant factors, and we might *as a compensatory strategy prohibit* the use of race and ethnicity just to ensure that they were not overused. (Schauer 2009: 187—emphasis added)

This concern that race might be "overused" is associated with what is known in the literature as the "exaggeration hypothesis" or the "kernel of truth" hypothesis, according to which stereotypes may have certain validity, but that perceptual processes are *likely to exaggerate these real differences between groups*.

Schauer's suggestion has two problems. First, even if the police were prone to exaggerate the difference between blacks and whites (and thereby unjustifiably disadvantage blacks), prohibiting the use of race would not be the only way to solve that

problem. An obvious and arguably better alternative would be to train the police to overcome their tendency to magnify the existing difference.

Second, and more importantly, Schauer's hypothesis that people will tend to "overuse" the category of race had already been empirically undermined thirty years before the publication of his (2009) book. In a widely cited empirical study of stereotypes the authors report their most important finding: "There is *not one instance* of significant stereotyping where the stereotype ratio is more extreme than the criterion ratio." (McCauley and Stitt 1978: 937—emphasis added).

In the same vein, a later review of the literature on the exaggeration hypothesis concludes that "the evidence available does not support the hypothesis of a large and consistent exaggeration effect in group stereotyping" and that in fact "stereotypes of stereotypes may be consistently exaggerated, perhaps more strongly and more consistently exaggerated than are stereotypes of real group differences" (McCauley 1995: 239–240 & 241). It is odd that Schauer doesn't mention this criticism of the exaggeration hypothesis given that he himself cites the very book in which this article appeared.

In this section I have argued that racial profiling is based on a perfectly good form of statistical reasoning. This doesn't make it automatically justified, all things considered, but it does make it rational (as a more efficient way of fighting crime than its alternative), as well as *prima facie* acceptable. Perhaps profiling has some negative aspects that outweigh its crime-fighting benefits and make it morally problematic. Possible, but it is far from clear that this is the case. The following section addresses this question.

3 Moral or Immoral?

Even if racial profiling is rational there can still be moral objections to it which can make it impermissible. Economist Tyler Cowen writes about racial profiling in a recent post on his widely read blog *Marginal Revolution* (January 11, 2017):

[S]tatistical discrimination, even if it is rational, does not excuse what are often crimes against innocent people. For instance, a man is far more likely to kill you than is a woman, but that fact does not excuse the shooting of an innocent man.

Yes, of course, the fact that X belongs to a group that has a higher proportion of killers does not, by itself, excuse killing X. On the other hand, killing a person who happens to be innocent can be excusable. After all, the right of self-defense does include a right to kill a suspected attacker even in situations where there is a non-zero probability that the killed person is innocent. In fact, this probability is *always* non-zero. The right of self-defense would simply cease to exist if it applied only to cases in which the probability of killing an innocent person would have to be literally zero.

Suppose now, *just for the sake of illustration*, that police officers are told they are allowed to draw a gun and be prepared to shoot if the suspect does not comply with their request only in cases in which the probability that they (or other innocent people) are in mortal danger is higher than, say, 0.5. What would this imply in the situations depicted in Figs. 4 and 5 above?

To paraphrase Hume, police officers should proportion their actions to the evidence at hand. Accordingly, they should be instructed to draw a gun and issue a warning when facing someone about whom all they know is that he is black and has characteristic E. The reason: the probability of a serious threat in that case would be over the prescribed threshold of 0.5.

Whites with characteristic E, however, would not, as such, fall under the protocol because in their case the level of threat would be too low (less than 0.1). But it by no means follows that whites should never be treated as a serious threat. Suppose that a white individual has characteristic E *and* additional threatening characteristic F, whereby the probability of someone having all three characteristics (W, E and F) being a serious threat is 0.6. Clearly the protocol would apply to such a person.

This all means, simply, that it would take *more specific evidence* for the protocol to kick in and apply to whites than to blacks. But—many will say—this is exactly what is wrong with racial profiling! Isn't it deeply unfair that two persons are treated *differently* by the police even when the *same* evidence is present in both cases?

Yes, a black person and white person with the same characteristic E will be treated differently. But the evidence is *not* the same in the two cases, because the evidence does not reduce to just E. The evidence in the black case is B & E, whereas in the other case it is W & E. As we saw in section 2, group membership matters. And Carnap's principle of total evidence requires that *everything relevant* be taken into account when calculating the probabilities that will guide action.

So, are blacks and whites treated *equally*? In an important sense, yes. For, independently of whether a person is black or white, everyone falls under the police protocol only if the level of suspicion in his case, *judging by all relevant evidence*, is above the threshold of 0.5. This applies equally to blacks and whites.

Assuming that the protocol is followed, compare the police treatment of those people with characteristics B & E, with the treatment of those with characteristics W & E & F. The level of suspicion in the two groups is the same: 0.6. The frequency of true positives (i.e. appropriate defensive actions against really dangerous suspects) will be the same for both groups (60%), and likewise for the frequency of false positives (40%). Generalizing, for any given level of suspicion (based on all relevant evidence), blacks and whites are treated the same by the police. Whatever the race, the question is always "What is the overall suspicion level in this particular case?", and after that the protocol is applied indiscriminately, one way or the other. Is this fair?

To answer that question, imagine that only blacks live in the U.S. and that the homicide statistics for the whole country correspond to what is presented in Fig. 4. Suppose that the police in this alternative U.S. have to determine the protocol specifying under what circumstances police officers would be allowed to draw a gun and issue the last warning to a person they see as a threat to themselves or others. Finally, suppose they agree that the permissible trigger for police action are situations that roughly correspond to the probability of at least 0.5 that a suspect will try to seriously injure or kill someone if not stopped. Consequently, officers are enjoined to take action when confronted with any person with characteristic E (because the key probability with respect to all such people would be 0.6).

The point of this thought experiment is to ask the following question: given that blacks in the alternative U.S. would settle on some threshold of threat as a justification for tough police action, why would they have a different view on this issue in the real

U.S. today? Put differently, if in the alternative U.S. the threshold of threat were set on the probability of 0.5—and this were done on the basis of all the main relevant goals (effective prevention of crime, protecting the life of police officers and other innocents, discouraging excessive aggressiveness of the police, etc.)—why would the mere additional existence of other groups, in the real U.S., be a reason to change the policy that was supposedly well grounded?

To repeat, the probability of 0.5 is introduced here completely arbitrarily and it plays no essential part in the argument. The question is this: whatever the approximate threshold of threat is agreed upon in the alternative U.S. as a justifiable reason for tough police action, why wouldn't this same, presumably sensible, threshold then also apply, *and with the same justification*, to the real U.S.? Or, again in slightly different terms: if blacks would accept that threshold in a country in which all inhabitants were black, why not also in a country where other groups also exist (and for whose members the same protocol applies)?

The whole point of introducing a threshold is to balance different considerations that are relevant and should decide about what kind of threat should legitimize an aggressive police response. Once the issue has been settled (based on the concern for public safety, protection of innocent lives, the right of police officers to defend themselves, etc.), it is not immediately clear why the decision should be revisited or reconsidered just because different groups in society happen to have different levels of criminality.

An analogy might be helpful here. Since different races have different frequencies of certain medical conditions, doctors sometimes take this into account and tailor their approach to patients based on race. Such a race-based approach to medicine has even been approved by the Food and Drug Administration (e.g. Saul 2005). Yet despite a possible uneasiness, no one seems to be upset about this. Also, the black community would probably not accuse doctors of racism or discrimination if they advised that blacks be scrutinized more closely for the possible presence of some medical problems more prevalent in their group. They would understand that this is proposed for their own good and that it would be downright silly if they were offended and demanded the same treatment as whites in the name of "fairness".

But the same logic applies to police work. Given that blacks have a higher incidence of criminality than whites and given that black criminals mainly target other blacks, shouldn't the police be allowed to keep this in mind and act on that true belief in some situations, if for no other reason than to *protect* law-abiding blacks?

Let us not forget that the issue of profiling is not raised only about race. In the case of sex, men are on average considerably more prone to violent behavior than women. Hence the police will sometimes (perhaps often) react *differently* to a man than to a woman although the two might behave in the *same* way (say, they both yell at officers, act erratically, fail to comply with officers' requests, reach for something in their pocket, etc.). Most people wouldn't see a problem if a police officer in this situation treated the man as a greater threat than the woman despite their apparently *identical* behavior. Prior probabilities matter.

Let us consider some of the moral arguments against racial profiling.

(a). *Identity*

Risse and Zeckhauser mention the problem that "individuals may feel resentment because they are profiled on a characteristic that partly constitutes their identity" (Risse

and Zeckhauser 2004: 145). But being a man also partly constitutes one's identity, and men rarely object to being profiled as men. So it seems that profiling cannot be wrong *only* because it uses a characteristic related to one's identity. There must be something more than this.

(b). *Group membership*

Another frequent objection to racial profiling, also mentioned by Risse and Zeckhauser, is that "people may feel resentment because they are treated in terms of a group membership, at the exclusion of their other characteristics, thus not as they deserve". Unfortunately, people cannot always be treated exactly as they deserve (whatever "deserve" might mean"). In tense situations in which police officers have neither the time nor the opportunity to acquire knowledge of the suspect's "other characteristics", they simply have to act on the little information that is available to them, which usually includes statistics about groups. It is a trade-off, and many would agree that, however undesirable it may be to treat someone partly on the basis of a statistical characteristic of the group to which the person belongs, if this sufficiently increases the probability of saving someone else's life, this approach can be justified. It is not always easy to decide where to draw the line between justified and unjustified group profiling, but still it would be very implausible to claim that group profiling is *never* justified.

(c). *Undermined trust*

A standard complaint is that racial profiling will undermine trust in the police among blacks. This fear would be justified if the police treated *differently* whites and blacks who presented the *same* level of threat after all relevant evidence was taken into account. Such behavior would indeed amount to irrational discrimination and could not be permissible. But what about the police merely using the group statistics (base rates) to estimate the posterior probabilities of danger in individual cases? If the different base rates (prior probabilities) are rooted not in prejudice but in reality, what would the complaint then be about?

Well, one could argue here that even if blacks were *wrong* in their belief that the police treat them unfairly, this *mistaken* opinion could still have two very real and highly undesirable effects: (a) blacks could lose respect for the entire criminal justice system, and (b) the continuation of racial profiling could (in principle) lead to even higher rates of black violent crime in the long term. Would this be a good reason to oppose racial profiling?

Not really. First off, the claim that X *could* happen is a weak argument for anything. The starting question should be: is it reasonable to expect that X *would* happen (because of racial profiling)? And, especially with respect to (b), initial skepticism seems advisable. For if a group received special attention from the police because of its comparatively high rate of violent crime and if at one point the police reduced their interest in suspicious activities of the group members, would you expect, as a result, the sharp *decline* of crime in that group? Maybe not. For, at least on the face of it, there would be a strong reason to expect exactly the opposite.

Suppose, however, that we establish to our satisfaction that in the situation as it is racial profiling would indeed be anticipated to bring about those highly undesirable effects, (a) and (b). What then? Terminate racial profiling? Again, not so fast.

Isn't the best way to deal with people's *wrong* beliefs to talk to them and convince them to accept *true* beliefs? Isn't this better than keeping them in the state of ignorance and adopting a condescending attitude toward them?

For the sake of concrete illustration, imagine that the police are good Bayesians paying attention to prior probabilities and that their interventions reduce black crime by 20%. But suppose further that blacks *wrongly* believe that the police interventions are driven by racial animus and, as a consequence, their resentment and anger drive up the black crime by 40%.

So we can picture the process as unfolding in three stages. At the beginning (with no police intervention) we can set the index of the black crime to be 1. In the second stage, the police intervention reduces crime by 20%, so the black crime rate goes down to 0.8. Finally, after blacks notice and misinterpret the police behavior their anger builds up and the crime rate increases by 40%, going up to 1.12.

The second stage is obviously good for society as a whole *and* for blacks themselves. What is the cause of the worst, third stage outcome? And how can it be avoided?

Admittedly, the first cause of the worst outcome is police interventions, which then cause black anger, which then causes the rise of black crime. On the other hand, it is clear that police interventions are beneficial in themselves and have a negative effect only indirectly, via the *mistaken* way they are perceived by blacks. Isn't it patently more reasonable here to try to explain and remove the mistake rather than remove the intrinsically beneficial causal factor itself?

Also, why wouldn't blacks be able to understand that? And why wouldn't they be able to appreciate the true explanation, abandon their wrong beliefs and realize that their anger was unjustified? There is something almost insulting in this attitude that a *false* belief of many blacks has to be treated as a given, as a fixed point or an unalterable fact, a harmful and self-destructive delusion that they will never be able to get rid of. What is the source of that supposed inevitability? Why would blacks be permanently incapable of grasping the simple truth about elementary probability relations and accepting the relevance of different base rates for the police dealings with violent crime? This attitude is so patronizing and contemptuous toward blacks that I would be tempted to call it "racist" if I didn't think that the inflated and indiscriminate use of this expression has long deprived it of any clear meaning.

Besides, the hypothesis that police profiling is caused by animosity toward blacks is undermined by repeated findings that white police officers are *not* more likely than black officers to kill black suspects (Lott Jr and Moody 2016: 14; Fachner and Carter 2015: 3), and that "black and white people shot and killed by the police are more likely than not killed by an officer of their own race" (Jan 2017). There is evidence that reactions of the police are not determined by their skin color but by their awareness of the world they live in.

It is well known that violent crimes committed by blacks are mostly directed at other blacks. For this reason any policy that would urge law enforcement officers to resort to tough intervention *less often* when dealing with black suspects would harm blacks themselves. How? In order to keep statistics about shooting incidents color blind, officers would have to behave *as if* the risk posed by black suspects is lower than they

think it really is (judging by all relevant information, including the base rates). But since such an inappropriate reaction would put their own life in great danger, policemen would often be reluctant to pursue apparently dangerous black criminals because this would place them between a rock and hard place: between responding appropriately to threat (for which they could be accused of racial profiling) or responding inappropriately (which would cripple their ability for self-defense). They cannot be blamed if they shun such lose-lose situations by simply avoiding confrontation.

However, failing to swiftly stop possibly armed and dangerous black suspects would endanger the lives of other, innocent blacks. The black community should therefore think hard about whether it is really in their interest to demand that the police be more restrained in encounters with black suspects given that this change in police behavior would inevitably lead to more innocent blacks being killed (including black children, who are often victims of drive-by shootings and gang wars).

(d). *The problem with innocents*

Recall that in the situation depicted in Figs. 4 and 5 the police would treat blacks with characteristic E as a greater threat than whites with the same characteristic (because of the huge difference between frequencies of homicide offenders among B&E and W&E). How would this affect the treatment of *innocent* people in each group? Since for all whites with characteristic E (and no other information) the level of risk they pose would be far below the critical threshold of 0.5 (which is necessary for a tough police response), the police would draw a gun on *no* innocent person in the W&E group. On the other hand, in the case of blacks with characteristic E (and no other information), since their level of risk is above the threshold, the police would resort to tough measures toward *all* innocent persons with characteristic E. So, looking at people with *the same specific evidence*, the police would react aggressively toward 30 black people who are in fact innocent (as they are all above the threat threshold of 0.5), and zero white people who are innocent (as they are all below the threshold).

So, merely being black rather than white would make an innocent person with E much more likely to be treated aggressively by the police. Isn't this a problem?

There are two responses to this question. First, there is a trade-off here. Since it is assumed that the police have to act, one way or the other, without being able to acquire more information and distinguish between the innocent and non-innocent, what should the policy be for dealing with such situations? If it is decided that the top priority is to make innocents in both groups (whites and blacks) *equally likely* to be treated aggressively by the police, then the fact of different priors should be disallowed. Fine. Suppose that this egalitarian policy is implemented by applying the base rates for whites to everyone. This means that the police would not react appropriately and on time to any of those blacks with characteristic E (60% of them, to be precise) who are an immediate mortal threat to police officers and others. Another thing to keep in mind is that many of the innocent people killed because of the new policy would themselves be black, e.g. black police officers. Is this additional loss of white and very often *black* lives a price worth paying for making the frequency of close encounters with the police equal for blacks and whites?

Second, so far I have been using the term *innocent* to refer to those blacks with characteristic E who do not really present a serious threat to anyone. This seems to suggest

that the tough way the police treats them is just something that happens to them and for which they bear no degree of blame at all. This increases the appearance of unfairness and sense of urgency to do something about it. But perhaps this impression is not entirely correct.

Characteristic E, which raises a red flag for the police and puts them in a self-defense mode, is typically associated with the person's behavior. The fact that the person belongs to an ethnic group that has a higher than average frequency of criminality is never by itself a sufficient reason for the police to prepare for an immediate violent confrontation. For this to happen, the person has, *in addition*, to *choose* to act in a way that law enforcement officers will then see as a cause for alarm and a reason to increase their level of alert. But obviously there is no inevitability about this; the person may well choose *not* to act that way. The decision is totally up to him.

So we see that a person stopped by the police is to a significant degree responsible for whether the situation will get out of hand and escalate into a scuffle that may lead to a loss of life. In the case of tragic outcome the person will not be without blame because he probably missed several opportunities to defuse the tension and end the conflict with no one being harmed.

A critical moment when any minimally reasonable person must realize that it is imperative to back down is after the police officer draws a gun and threatens to shoot. If one continues to be combative even after that dire warning and in the end gets killed, one is at least partly responsible for one's own death (even if one did not intend to harm anyone!).

All this points to a possible way of reducing the number of blacks killed by the police, namely by explaining to them the utmost importance of restraint in their encounters with law enforcement officers in precarious circumstances. If they realize that they are in fact largely in control of those situations, in the sense that it is highly unlikely that a police officer would shoot a fully cooperating suspect, they should stop behaving in a way that could be perceived as threatening. This will help minimize the danger of a bad outcome. There will always be time later for lodging a complaint against the police if the officers violated their code of conduct.

Also, if there is a perception that the police are more on edge when dealing with blacks than with whites, an effort should be made to allow for the possibility that this is *not* the result of their racism or hatred toward blacks but is due to the established fact that, statistically speaking, blacks *are* more dangerous than whites. Since law enforcement officers are not robots stripped of all emotions, when they fear for their lives they cannot help but react differently to really existing different levels of danger.

If blacks, however, accept a rival narrative according to which the police are trigger-happy toward blacks mainly because of the officers' irrational hatred and bigotry, then an idea of armed self-defense or even "pre-emptive" attacks on the police (which have already happened) may seem to make some sense. In my opinion, there is no good evidence to support this alternative narrative, but it has been advocated by many supposed friends of the black community. If the hypothesis is false (as I think it is), its further spread will needlessly cost many future lives on both sides.

(e). *More knowledge, less action*

One of the most bizarre arguments against racial profiling is defended by the well-known philosopher Philip Kitcher (at Columbia University), who argues that a wider context is highly relevant:

This is most evident in relation to crime, where "profiling" is defended by appeal to statistics, without any consideration for background circumstances that might lie behind the numbers cited. To inaugurate or continue a policy of stopping young men of a particular race, without doing additional experiments to see what might be done to increase the opportunities and prospects of members of the targeted group, is to conduct an illegitimate experiment, one that an informed, mutually engaged, comprehensive deliberation would almost certainly reject. (Kitcher 2016: 252)

There are several problems with this argument. First, there is a clear division of labor between the police (who fight crime) and criminologists (who try to understand why some people or groups engage in crime). Why should these two tasks be linked in the way Kitcher suggests, so that the work of the police should depend on the efforts of criminologists? Don't the police operate in the here and now, without having the time or luxury to wait for social scientists to sort out the correct explanation of different frequencies of criminality in different groups and propose remedies? Or, to put it briefly, if you think somebody is about to attack you with the intention to murder you, your immediate concern will not (and should not) be to try to figure out why he and other aggressive people from his group became that way.

Second, Kitcher assumes, without offering any evidence in support of his assumption, that the higher level of crime among blacks is due to their lack of "opportunities and prospects". I am not saying the assumption is false. My objection is merely that it should not be accepted without doing serious research and considering alternative hypotheses. Some scholars believe that in the search for causes of the rampant black crime we should look, among other things, into the strand of black culture that glorifies violence (e.g. rap music), contempt for those who try to acquire skills and education (who are accused of "acting white"), disintegration of black family, the bad influence of "redneck culture" (Sowell 2005), opposition to the police from within black communities, etc. These and other factors may well be an important part of the story, and in that case turning a blind eye to them (as Kitcher does) would only do disservice to blacks.

Third, when Kitcher prophesies that "an informed, mutually engaged, comprehensive deliberation would almost certainly reject [racial profiling]", this is a gratuitous assertion. The idea that it is possible to predict both a definitive answer to a currently intractable question of such enormous complexity and also the practical implications of that answer is ludicrous. Furthermore, one should notice how suspiciously the results of future research that Kitcher anticipates happen to coincide so neatly with the views of the political left, with which he himself identifies.

(f). *Using race in criminal trials is wrong, therefore...*

A reader of an earlier draft of this article raised the following question:

Most of us would find something immoral about using race to determine a verdict of criminal guilt or innocence in a trial when the evidence is less than dispositive, which it almost always is. Perhaps we shouldn't be so squeamish, and our

misgivings are irrational, but they are certainly real, and if they are valid, would they also apply to police decisions?

It seems to me that there are at least three important and morally relevant differences between the two situations. First, in the police case (but *not* in the trial case) there is typically a particular person—a police officer or an innocent bystander—who appears to be in grave danger and whose life will be put at much higher risk if the suspected attacker’s race is not taken into account when deciding how to respond. So, *only* in the police scenario the price for non-profiling would be very high, specific, and immediate. Second, in the trial case (but *not* in the police case) there is always enough time to collect a lot of other evidence that would soon swamp the importance of race, decrease its probative value and consequently make it easier to reasonably ignore. And third, in contrast to the trial case where innocent subjects of racial profiling would *not* have a foolproof way of avoiding the worst possible outcome (i.e., a harsh unjust sentence, including possibly a death penalty), in the police case both innocent and non-innocent subjects of racial profiling can in most situations easily avoid the worst possible outcome (i.e., being shot by the police) by simply complying with police requests and not behaving in a threatening way.

All these three differences indicate that it is much harder to justify racial profiling in the trial case than in the police case. Therefore, as long as the alleged similarity between the two kinds of racial profiling is not established in a more convincing way, we should not infer that police racial profiling is morally impermissible merely on the grounds that racial profiling in criminal trials is regarded as morally unacceptable. An argument connecting the two is needed to legitimate that step.

4 Conclusion

What is the take away lesson of this article? The first part shows that, contrary to widespread opinion, racial profiling by the police is in principle rational. The second part argues that the main moral arguments against racial profiling are not convincing.

Many will find this message disheartening. What are we then to say to those who still want to keep alive that dream of a world in which people will *never* be judged by the color of their skin, but only by the content of their character? One might respond as follows: “By all means continue to dream about that wonderful possible world. But when you are back on planet Earth try to understand that people will keep adjusting their behavior to the reality of those still existing group differences, especially when ignoring them would put their own life or lives of others in great danger.”

Acknowledgements I am grateful to Nathan Cofnas, Rafael De Clercq, Steven Pinker, Ante Sesardić, Matej Sušnik and Omri Tal for useful comments on earlier drafts. None of them, however, should be assumed to agree with the main claims of this paper.

References

Balko, R. (2016). Why It’s Impossible to Calculate the Percentage of Police Shootings That Are Legitimate, *Washington Post*.

- Beaver, K.M. et al. (2013). No Evidence of Racial Discrimination in Criminal Justice Processing: Results from the National Longitudinal Study of Adolescent Health, *Personality and Individual Differences*, 55, 29–34.
- DOJ (2003) Guidance Regarding the Use of Race by Federal Law Enforcement Agencies.
- DOJ (2011) Homicide Trends in the United States, 1980–2008.
- DOJ (2014) Guidance for Federal Law Enforcement Agencies Regarding the Use of Race, Ethnicity, Gender, National Origin, Religion, Sexual Orientation, or Gender Identity.
- Gelman, A. (2016). About that Claim that Police are Less Likely to Shoot Blacks than Whites, *Statistical Modeling, Causal Inference, and Social Science*. July 14 (<http://andrewgelman.com/2016/07/14/about-that-claim-that-police-are-less-likely-to-shootblacks-than-whites/>)
- Fachner, G. & Carter, S (2015). *An Assessment of Deadly Force in the Philadelphia Police Department*, Community Oriented Policing Services, U.S. Department of Justice.
- Fryer, R. G. Jr, (2016). An Empirical Analysis of Racial Differences in Police Use of Force, working paper 22399.
- Jan, T. (2017). Study: Hiring More Black Cops Won't Stop Fatal Police Shootings of Black Citizens, *Washington Post*.
- Kitcher, P. (2016). Reply to Cartwright and Marcellesi. In M. Couch & J. Pfeifer (Eds.), *The Philosophy of Philip Kitcher*. New York: Oxford University Press.
- LaCapria, K. (2016). A 'Harvard Study' Doesn't Disprove Racial Bias in Officer-Involved Shootings, *Snopes*.
- Lind, D. (2016). Here's Why I'm Skeptical of Roland Fryer's New, Much-hyped Study on Police Shootings, *Vox*.
- Lott, J. R. Jr. & Moody, C. E. (2016). Do White Police Officers Unfairly Target Black Suspects? (unpublished).
- Mac Donald, H. (2003). *Are Cops Racist?* Chicago: Ivan R. Dee.
- Mac Donald, H. (2016a). *The War on Cops*. New York: Encounter Books.
- Mac Donald, H. (2016b) Are the Police Racist, Prager U YouTube video.
- Mac Donald, H. (2016c). Academic Research on Police Shootings and Race, *Washington Post*.
- McCauley, C. R. (1995): Are Stereotypes Exaggerated? A Sampling of Racial, Gender, Academic, Occupational, and Political Stereotypes, in C. R. McCauley et al. (eds.). *Stereotype Accuracy: Toward Appreciating Group Differences*: American Psychological Association. 215–243
- McCauley, C., & Stitt, C. L. (1978). An Individual and Quantitative Measure of Stereotypes. *Personality Social Psychology*, 36, 1978, 929-940
- Risse, M., & Zeckhauser, R. (2004). Racial Profiling. *Philosophy & Public Affairs*, 32, 131–170.
- Saul, S. (2005). F.D.A. Approves a Heart Drug For African-Americans, *New York Times*.
- Schauer, F. (2009). *Profiles, Probabilities, and Stereotypes*. Cambridge: Harvard University Press.
- Simonsohn, U. (2016). Teenagers in Bikinis: Interpreting Police-Shooting Data, *Data Colada*.
- Sowell, T. (2005). *Black Rednecks and White Liberals*. New York: Encounter Books.
- Tetlock, P. E., et al. (2000). The Psychology of the Unthinkable. *Journal of Personality and Social Psychology*, 78, 853–870.
- Tversky, A., & Kahneman, D. (1982). Evidential Impact of Base Rates. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge: Cambridge University Press.