

On Selective Emphasis, Broad Agreement, and Future Directions: Reply to Ross, Winterhalder, & McElreath

Joseph Cesario
Michigan State University

Ross, Winterhalder, and McElreath (in press) have provided a valuable and important piece of work in the context of addressing an earlier paper by Cesario, Johnson, and Terrill (2019). Cesario et al. argued that population proportions are a misleading benchmark for testing racial disparities in deadly force and that violent crime rates are instead a more appropriate benchmark. Using such rates, Cesario et al. reported two major findings: (1) no evidence of anti-Black disparity when looking at all fatal shootings, and (2) uncertainty regarding shootings of unarmed citizens due to the small number of such cases.

In this reply, I address three aspects of Ross et al. (in press)'s paper: (1) that the results presented in Ross et al. in fact *replicate the central findings of Cesario et al.*; (2) that Ross et al. provide a very important corrective and approach, but one that has its own assumptions and which may have overestimated anti-Black disparity in shootings of unarmed, non-aggressing civilians; and (3) that there are productive

Joseph Cesario is professor of psychology at Michigan State University.

For reasons that remain unclear to me, *SPPS* peer-reviewed and published Ross et al.'s manuscript without asking me to review the submission. Days before Ross et al.'s manuscript was to appear published online, the AE first notified me of this paper and invited me to write a response. My response went through two rounds of peer review, during which Ross himself recommended publication of my response and stated the following during the review process: **"I agree, more or less, with all of the substantial claims made here."**

The AE ultimately rejected publishing my response "for two primary reasons." First, that it was not "responsive enough" to Ross et al.'s paper and that it "underplays the importance of [Ross et al.'s] findings." (I later pointed out to the AE that I praised Ross et al.'s work *more than a dozen times* in my short reply.) Second, that the reply was not sufficiently "relevant to the focus of this particular journal and its readership," and that the AE "had hoped that your reply would make stronger connections to social and personality science." This, despite the fact that the original request by the AE was to provide my "perspective on the recently accepted manuscript and where you would stand with your original findings given this new work" and to "lay out a path for moving forward," both of which I did in my reply.

My rejected submission to *SPPS* is reproduced here.

paths forward to be informed by both approaches.

To provide the short summary up front: **for 90-95% of fatal shootings**, both Cesario et al. (2019) and Ross et al. (in press) agree that there is no evidence of anti-Black disparity and instead report substantial evidence of anti-White disparity.

For the remaining 5-10% (shootings of unarmed and non-aggressing civilians), Cesario et al. claimed the data were too uncertain but Ross et al. claim positive evidence of anti-Black disparity. Regarding Cesario et al.'s claim, Ross et al. demonstrate that the method used by Cesario et al. can mask actual racial disparity in these shootings. Rather than being "too uncertain," the method used by Cesario et al. may have introduced an opposing bias that masked real anti-Black disparities. Regarding Ross et al.'s claim, in this comment I argue that the specific model used by Ross et al. may have overestimated anti-Black disparity in shootings of unarmed, non-aggressing citizens. While this discrepancy cannot be resolved in this comment, Ross et al. have provided a method that can be productively used to yield answers to this question.

Finally, both Cesario et al. and Ross et al. agree that comparisons of the outcome (being fatally shot) to population proportions are misleading.

Selective Emphasis

The reader could be forgiven for coming away with the impression that Ross et al. (in press) undermines or contradicts the main conclusions of Cesario et al. (2019). This is because the key statement:

we recover a principle finding of Cesario et al. (2019): racial disparities in the killing of armed suspects by police are proportional to the relative rates of violent criminality

first appears buried halfway through the manuscript, with the title, abstract, and most of the document instead stressing the differences in findings between Ross et al. and Cesario et al.

But what exactly are the differences and similarities? Starting with the main difference, Ross et al. (in press) claim that the data reveal "strong and statistically reliable evidence of anti-black racial disparities in the killing of unarmed

Americans."¹ This differs from but does not strongly contradict the conclusions of Cesario et al. (2019), because Cesario et al. were *agnostic* with respect to shootings of unarmed citizens, stating in the abstract that for unarmed shootings “data are too uncertain to be conclusive.”

We took this position because of the small number of cases on which these estimates were based: only about 20 Black citizens and 30 White citizens per year. Given the high uncertainty around these estimates, we believed that the most reasonable position was to be clear about the uncertainty and conclude that the data were too sparse to draw any conclusions one way or another. Ross et al. (in press) however demonstrate that the benchmarking approach of Cesario et al. (2019) may have introduced a statistical bias that masked anti-Black disparity in these shootings and they believe they have enough certainty to make a positive claim. Indeed, they do have credible posterior probability estimates with a 90% interval.

Thus the main difference in results between Cesario et al. (2019) and Ross et al. (in press) is that, for killings of unarmed, non-aggressive citizens (again, roughly 5-10% of shootings), we expressed uncertainty and did not take a strong position whereas Ross et al. conclude anti-Black disparity.

As for the similarities, Ross et al. (in press) confirm that there is no anti-Black disparity in shootings of armed or aggressing citizens (the far majority shootings) once violent crime rates are considered. Indeed, they confirm not just a lack of anti-Black disparity but reveal an anti-White disparity in most of their calculations (see their *Fig. 2*). That is, *the most certain results in all of Ross et al.’s analyses confirm that there is no anti-Black disparity in fatal police shootings of armed or aggressing citizens.*

Therefore, even without addressing anything about the validity of Ross et al. (in press)’s approach, when considering the overwhelming cases of fatal police shootings, *accounting for violent crime rates eliminates anti-Black disparities.* Nothing in Ross et al.’s paper contradicts this point.

The next question is, what is the quality of the approach used by Ross et al. (in press)?

Ross et al.’s Approach

Cesario et al. (2019) compared the rate of fatal shootings for Black and White citizens against the rate of violent crime for Black and White citizens. We computed odds ratios in this way rather than the common method of comparing rates of fatal shootings to the population proportions of all Black and White citizens in the U.S. When benchmarking against violent crime rates, Cesario et al. found no evidence of anti-Black disparity in fatal police shootings.

Cesario et al. (2019) failed to explicitly describe the underlying causal model and to produce estimates of disparities based on such a causal model. Ross et al. (in press) work out

the causal model implicit in Cesario et al.’s work and show that simple comparisons against violent crime rates can, under reasonable conditions, mask anti-Black disparity.

Ross et al. (in press) then apply their model to data similar to that used in Cesario et al. (2019) and show two findings: no anti-Black disparity for shootings of armed or aggressing citizens (roughly 90-95% of all shootings) and credible anti-Black disparity among the remaining shootings of unarmed, non-aggressing citizens.

To evaluate the contribution of Ross et al. (in press), it is useful to separate their quantitative approach in general from the specific model they implement and the assumptions inherent in that model.

As a general approach, Ross et al. (in press) should be commended for introducing a formal causal model into this research tradition. They highlighted the failure of Cesario et al. (2019) in describing a causal model and showed how this could produce misleading results, particularly in the case where Cesario et al. benchmarked fatal shootings of unarmed, nonaggressing citizens against violent crime rates. Beyond Cesario et al., many researchers have used simple percentage comparisons in understanding racial disparities (e.g., Goff, Lloyd, Geller, Raphael, & Glaser, 2016) and so Ross et al.’s contribution is widely applicable. This is a very important advance, and (as I address in more detail below) it provides a model moving forward for more complete and detailed analyses of the full process by which police-citizen interactions move from initial encounters to ending in fatal shootings.

Regarding the specific instantiation of the approach introduced by Ross et al. (in press), the major concern is the division of the population into “criminal” and “non-criminal” and the use of the *ratio of non-criminality* as a multiplicative correction to produce an “unbiased” estimate. Specifically, Ross et al. argue that although violent crime ratios are the relevant corrective for producing unbiased estimates of killing “armed criminals,” the *ratio of non-criminality* is the relevant corrective for producing unbiased estimates of killing “unarmed non-criminals.”

However, their value, $\frac{1-\alpha_W}{1-\alpha_B}$, is merely the entire non-violent race-specific population, that is, everyone who is not a violent criminal. The problem with this corrective is that it fails to account for the fact that among nonviolent citizens, encountering police is still a prerequisite for being fatally shot. If encounter rates differ between Blacks and Whites for the nonviolent population, then the estimates produced by Ross et al. will not be unbiased. Indeed, if encounter rates

¹“Unarmed” here technically refers to “Unarmed and not aggressing” because citizens can pose a deadly threat even when unarmed. The division concerns citizens who are shot while “armed and/or aggressing” versus citizens who are shot while “unarmed and not aggressing.” Ross et al. (in press) sometimes use “unarmed” and sometimes use “unarmed and not aggressing.”

are *higher* for Black nonviolent citizens, then the analyses will *overestimate* anti-Black disparity in unarmed shootings by biasing the size of the nonviolent citizen population to appear smaller for the Black population than it actually is.²

In other words, on the "non-criminal" side, Ross et al. (in press) return to the default position that Cesario et al. (2019) were arguing against: that encounter differences between Black and White Americans can be ignored when computing the likelihood of being fatally shot by the police. Cesario et al. took issue with the general *per capita* population comparisons that are typically used, but Ross et al.'s model has the same problem on the "non-criminal" calculations: treating the entire non-criminal population as a monolithic entity that is equivalent in all relevant ways between Black and White citizens.

Additionally, the simplicity of the specific model used by Ross et al. (in press) obscures the reasons *why* anti-Black disparity might exist among unarmed, non-aggressing citizens. To be perfectly clear, this is *not* a fault of Ross et al. as they are very precise in referring to anti-Black *disparity* rather than police officer *bias* throughout their manuscript; thus I am not faulting them for this. I am merely raising the issue as a means of suggesting productive future directions, to which I now turn.

Directions Forward for a Deeper Understanding

Beyond the specific problem of potentially overestimating anti-Black disparity in shootings of unarmed, non-aggressing citizens, a broader problem is that merely observing an anti-Black disparity does not tell us why the disparity exists (see also Tregle, Nix, & Alpert, 2019).

Given the specific model used by Ross et al. (in press), any of the following could produce the observed anti-Black disparity:

- Racial bias on the part of police officers in discretionary stopping at the initial encounter stage for non-violent, unarmed citizens, coupled with no racial bias on the part of police in the decision to shoot
- Racial disparities in citizens' calls for service at the initial encounter stage for non-violent, unarmed citizens, coupled with no racial bias on the part of police officers in the decision to make contact *or* in the decision to shoot
- Police officers applying a race-neutral decision rule to make contact with citizens combined with a constant race-neutral error rate in the decision to shoot, with no racial bias anywhere in the decision process

On this last point, it is possible for Ross et al. (in press)'s model to generate anti-Black disparity in shootings of unarmed, non-aggressing citizens while police officers are making completely race-neutral decisions at every stage. Offi-

cers, like all humans, will make some number of errors in their decisions. We can define shootings of unarmed, non-aggressing citizens as errors (as in the cases of Philandro Castille or Dylan Noble).³ Suppose that officers make errors in a completely race-neutral way, e.g., they mistakenly perceive harmless objects as weapons at a constant rate for both Black and White citizens. Suppose also that officers make initial contact with non-aggressing citizens by applying a race-neutral decision rule in a completely race-blind manner (e.g., pulling over citizens with broken taillights). If the criterion for initial contact varies systematically with race (e.g., if lower income citizens are more likely to drive with broken taillights and Black citizens are more likely to have lower incomes), then the model will generate anti-Black disparity in shootings of unarmed, non-aggressing citizens *while race plays no role in any stage of the decision process*.

Therefore although the finding of anti-Black disparities in shootings of unarmed, non-aggressing citizens is *consistent* with data from experimental social psychology showing racial biases in various types of perceptions (e.g., Wilson, Hugenberg, & Rule, 2017), it is also consistent with non-bias explanations or explanations in which bias enters at an earlier encounter stage and not at the stage of the decision to shoot. Moreover, other data in experimental social psychology provides evidence that strong situational information such as police dispatch information can overwhelm racial bias in the decision to shoot (e.g., Johnson, Cesario, & Pleskac, 2018), supporting a non-racial bias account of the anti-Black disparity in this type of fatal police shooting.

Given this, one path forward is to develop more expansive decision-tree models that estimate probabilities running the complete course of events, from the entire population, through different types of initial contact with citizens, through different types of shootings. This has the potential to isolate different possible sources of police officer bias in explaining different racial disparities at different stages of the encounter process. Moreover, more complex models can incorporate various citizen characteristics, such as the likelihood of violating laws that allow police to make discretionary contact or the likelihood of escalating initially non-violent interactions with police to the level of violence, if such data are available.

Hence what is needed is, as always, more and better policing data. If more complete encounter data were available, various proxies for encounters would not be needed. Both

²I am using the phrase "nonviolent citizens" to refer to those citizens who have contact with police but are not captured in the counts of "violent criminals." However, technically "nonviolent citizens" should be "citizens who break nonviolent laws or who in some other way attract the attention of law enforcement".

³It is worth noting that, relative to the total number of police-citizen contacts per year, officers make such errors in roughly about .0001% of cases.

Cesario et al. (2019) and Ross et al. (in press) are limited in using different types of report data (including self-reported victimization, homicide data, etc.) to estimate violent and non-violent encounters with the police. Of course these are not perfect. Moreover, all the analyses discussed herein are at the national level, and there is significant department-to-department variation in important policing procedures and outcomes (Ross, 2015). At the same time, however, merely pointing out that all data are imperfect approximations does not necessarily mean there are systematic errors that undermine any contribution. This further highlights the importance of drawing conclusions from as many data sources as possible, as Cesario et al. did with their multiverse analysis.

Conclusion

Cesario et al. (2019) showed that fatal police shootings benchmarked on violent crime rates fail to show the anti-Black disparity present when benchmarked against population levels. We also noted the high levels of uncertainty in shootings of unarmed, non-aggressing civilians.

In their critique, Ross et al. (in press) note that Cesario et al. (2019) failed to explicitly specify the causal model behind their analysis. Ross et al. specify such a model and in doing so they produce an important advance in calculating racial disparities. Moreover, they highlight how the application of the specific benchmarking of Cesario et al. was in error when applied to shootings of unarmed, non-aggressive citizens.

Ross et al. (in press)'s advance produces two main findings: (1) for nearly all police shootings, there is no evidence of anti-Black disparity; (2) for shootings of unarmed and non-aggressing citizens, there is credible evidence of anti-

Black disparity. The general approach by Ross et al. is very important and can produce productive further research. In this comment, I bring out some of the assumptions of the specific model applied by Ross et al. with the hopes of generating some of these further applications.

References

- Cesario, J., Johnson, D. J., & Terrill, W. (2019). Is there evidence of racial disparity in police use of deadly force? Analyses of officer-involved fatal shootings in 2015–2016. *Social Psychological Personality Science*, *10*, 586–595.
- Goff, P. A., Lloyd, T., Geller, A., Raphael, S., & Glaser, J. (2016). *The science of justice: Race, arrests, and police use of force*. Center for Policing Equity.
- Johnson, D. J., Cesario, J., & Pleskac, T. J. (2018). How prior information and police experience impact decisions to shoot. *Journal of personality and social psychology*, *115*(4), 601.
- Ross, C. T. (2015). A multi-level bayesian analysis of racial bias in police shootings at the county-level in the United States, 2011–2014. *PLOS One*, *10*(11), e0141854.
- Ross, C. T., Winterhalder, B., & McElreath, R. (in press). Racial disparities in police use of deadly force against unarmed individuals persist after appropriately benchmarking shooting data on violent crime rates. *Social Psychological and Personality Science*.
- Tregle, B., Nix, J., & Alpert, G. P. (2019). Disparity does not mean bias: Making sense of observed racial disparities in fatal officer-involved shootings with multiple benchmarks. *J. Crime Justice*, *42*(1), 18–31.
- Wilson, J. P., Hugenberg, K., & Rule, N. O. (2017). Racial bias in judgments of physical size and formidability: From size to threat. *Journal of personality and social psychology*, *113*(1), 59.