**No Evidence of Bias When Using Inappropriate Test for Bias: Comment on Cesario,
Johnson, & Terrill 2018**

William H. Ryan¹ & Ellen R. K. Evers¹

[1] Department of Marketing

UC Berkeley, Haas School of Business

Corresponding author: William Ryan (wryan@berkeley.edu)

Posting date: 2/11/2020

Note from the authors:

For context, this commentary was rejected from Social Psychological and Personality Science, the journal which published the original article, for two reasons. First, another comment on the Cesario et al. article reanalyzing this data with better benchmarks (and reaching opposite conclusions) was submitted – we are curious to see this paper, and slightly concerned because the issues we point out here are difficult to circumvent. Second, that it is a matter of opinion which question is the right one to ask, and that our argument was just that Cesario et al. had asked the wrong research question. This reading of our comment is incorrect, and not what we are arguing at all. Asking different research questions is fine, but we argue that Cesario et al.’s paper has a fundamental logical/methodological flaw which means that they are unable to answer their own research question.

Abstract

A recently published paper claimed that there is no evidence of anti-black bias in police shootings (Cesario, Johnson & Terrill 2018). They compare the ratio of blacks and whites shot by police to the number of violent crimes committed by each race, instead of to population levels. Based on this they find no evidence for anti-black bias in police shootings, actually finding an anti-white bias in many cases. However, the methods used by the authors are unable to determine if bias exists. We illustrate by showing that their method is unable to distinguish between a hypothetical world in which no bias exists, and one in which severe anti-black bias exists. Analysis of the same shootings database used by the authors further indicate the authors' assumptions do not match with the real world. Given this, any conclusions being made from the data and analysis used in the paper are unwarranted.

Keywords: statistical methods, police use of force, officer involved shootings, race bias

NO EVIDENCE OF BIAS WITH INAPPROPRIATE TEST

Recently a paper was published claiming that there is no evidence of bias against black people in police shootings, and if anything, bias against white people appears more likely (Cesario, Johnson &, Terrill 2018). The authors argue that rather than look at population-level data, the number of people shot by race should be interpreted in light of possible police interactions. Therefore, they compared the ratio of blacks and whites shot by police to the number of violent crimes committed by each race. Using this strategy to analyze the data, they claim that not only is there no anti-black bias in police shootings, there is actually a severe anti-white bias in many cases, with some benchmarks suggesting the odds of whites being shot by police are 3.9 times higher than blacks. We do not have the data to better answer the question the authors try to answer in their paper, rather, we will show that the methods used in this paper are invalid and cannot answer the key question the authors are trying to answer.

One may wonder, isn't an answer arrived at through an imperfect method better than having no knowledge at all? As we will show, the method used in this paper can lead to incorrect conclusions, and it's better to know you don't have an answer to a question than to believe in a dramatically wrong one. We will illustrate why this is the case with two hypothetical worlds; one in which no bias exists, and one in which severe anti-black bias exists. We will show that the author's methods cannot distinguish between these two worlds and in both cases will claim there is no bias. Given this, any conclusions made from the analysis and data in the paper are completely unwarranted.

Effectively, by using *violent* crime as a benchmark, the authors are assuming that only *violent* criminals are shot by the police. In reality, there are many situations in which people can

NO EVIDENCE OF BIAS WITH INAPPROPRIATE TEST

interact with the police. Police stop violent and non-violent crime suspects, conduct discretionary stops, and carry out welfare checks requested by family members. All of these are situations in which people interact with the police, and where there is the potential for people to be shot.

Crucially, the authors' argument completely relies on the assumption that police *only* or *almost only* shoot violent criminals. If police regularly stop and shoot individuals who were not suspected of violent crimes when they were stopped, then violent crimes are not a viable proxy for police encounters. This is because while the absolute number of violent black and white criminals may be roughly the same, the absolute number of non-violent black criminals and non-criminal black people is much lower than the absolute number of white non-violent criminals and non-criminals.

To be clear, we are not arguing here that Cesario made an error in choosing the denominator or base rate for his benchmark, we are arguing that his numerator (all people shot by police) and his denominator (violent criminals), don't match up. We can illustrate this idea with an intuitive analogous example: If you want to determine if there is gender bias in NSF winners in general, you should compare the gender ratio of NSF winners to the gender ratio of applicants as a whole. If you instead want to determine if there is gender bias in NSF winners in Engineering, you should compare the gender ratio of NSF winners in Engineering to the gender ratio of applicants within Engineering. What you should absolutely not do is compare the gender ratio of NSF winners *overall* to the gender ratio of applicants within *only* Engineering, which, just to illustrate, would show a massive "anti-male" bias. This incorrect approach is the method Cesario uses in his paper – he compares all those who are shot by police to only violent

NO EVIDENCE OF BIAS WITH INAPPROPRIATE TEST

criminals. Cesario's method could only be correct if and only if every person shot by police was a violent criminal.

To illustrate why this difference matters, we will next present two hypothetical worlds, one in which there is no bias and the police only shoots violent criminals, and one in which there is massive anti-black bias in the shooting of non-violent criminals. For both of these hypothetical worlds, the method the authors used would conclude no evidence of bias, despite one of them exhibiting massive anti-black bias. After this example, we present evidence that under the most conservative estimates, 40% of black people recently shot by police were shot during a stop in which they were not suspected of committing a violent crime, challenging the validity of the author's assumptions.

The author's method can't distinguish between bias and no bias.

To illustrate why it matters whether the police only shoot violent criminals, consider two hypothetical worlds: World A and World B. World A and World B are different in only one way: in World A, the police only shoot violent criminals, while in World B, they shoot both violent and non-violent criminals. We choose numbers in this example to roughly match reality, but the logic of the comment applies for a wide range of specific assumptions. In both world A and world B, there are 4,000 white people and 1,000 black people, about the same population ratios as in the US. Like the argument of the paper, in both World A and World B, black people are four times more likely to commit a violent offence. Like in reality, in world A and world B black people are, on the population level, more likely to get shot. Approximating reality, assume that

NO EVIDENCE OF BIAS WITH INAPPROPRIATE TEST

black people are, compared to population levels, four times more likely to get shot, and that therefore an equal absolute number of blacks and whites get shot and killed, let's assume 40 each. Finally, assume that both about a quarter of white and black people interact with the police, and that these interactions are primarily for non-violent crimes—only 1% of white interactions are for violent crimes, and 4% of black interactions, resulting in 100 violent crime interactions for each group¹.

If we consider victims by violent crime-numbers in World A, there appears to be no bias. Of the 100 white violent criminals, 40 get shot, meaning 40% of white violent criminals get shot. Of the 100 black violent criminals, 40 get shot, meaning 40% of black violent criminals get shot. As a result, conditional on being a violent criminal, there is no evidence for racial bias in police shootings.

World A essentially represents the analysis and conclusions presented in the paper.

Now, let's look at World B. World B is very similar to World A, the only difference is the type of criminal that gets shot by police. World B has the same populations, number of shooting victims, and number and types of interactions by race as World A. **However, in World B, only 60% of the victims of police shootings are suspects of violent crimes.**

Now the numbers for calculations of bias are different:

---

¹ This is a conservative assumption, if black people are less likely to call the police for non-violent crimes; the following examples would, if anything, be more extreme (for a discussion of why this may occur, see Desmond et al. 2016)

NO EVIDENCE OF BIAS WITH INAPPROPRIATE TEST

Among white people; 24 out of 40 victims come from the 100 white people interacting with police due to violent crime resulting in, for white people, 24% of violent white criminals getting shot.

Among black people, 24 out of 40 victims come from the 100 black people interacting with police due to violent crime resulting in, for black people, 24% of violent black criminals getting shot.

So, even in World B there is no evidence of bias when only violent crime is considered.

However, for the nonviolent police interactions, the following pattern emerges:

For white people; 16/40 victims come from the 1,000 police-interactions that do not involve suspicion of a violent crime, resulting in a $16/1,000 = 1.6\%$ chance of getting shot if you're a non-violent criminal white person.

For black people; 16/40 victims come from the 250 police-interactions that do not involve suspicion of a violent crime, resulting in a $16/250 = 6.4\%$ chance of getting shot if you're a non-violent criminal black person².

The analysis in the paper, applied to world B, would only compare 40 (white victims) divided by 100 (white violent criminals) with 40 (black victims) divided by 100 (black violent criminals) and conclude there was no bias, while in reality black non-violent criminals are over six times more likely to get shot.

Clearly, the entire paper stands and falls with the assumption that the police *only* shoot violent criminals.

---

² Here we assume that the police are only shooting criminals, however this point still holds even if the police also shoot non-criminals, as the population of black people is lower than the population of white people.

Do police only shoot violent criminals?

In the paper, the authors argue that it is reasonable to assume the police only shoot violent criminals, but do not really back up this claim. Instead the authors argue that most shootings happen *in situations of violence*, and argue that the police rarely shoot people who aren't criminals *of some kind* by the *time of the shooting*. This is problematic. *By definition*, as soon as people are getting shot, situations are violent, the crucial question is *when* the situation became violent.

For example, imagine relatives ask for a welfare check on an elderly man who is recovering from surgery, but police are unable to contact him. Eventually they break into his home at midnight. He investigates the break in while holding a gun, and when officers see him they immediately shoot and kill him³. While the shooting itself was violent, the cause of the police being on the scene was not. It is possible, even likely, that police interactions with people who are either not criminals at all or are non-violent criminals turn violent when the police arrive at the scene, ultimately leading to a shooting. But it is important to distinguish between violent crime leading to a police response, and a police response leading to violence. If stops do not begin due to a violent crime, then violent crime statistics obviously cannot be used as a proxy for the number of stops.

---

³ While this example is hypothetical, it is very similar to the circumstances in which James Howard Allen was shot: <https://wapo.st/31K747C>

NO EVIDENCE OF BIAS WITH INAPPROPRIATE TEST

While it is difficult to get exact data here, it's clear that for the author's conclusion to hold a high benchmark is required. For their analysis to make any sense, they do not require *some* of the fatal police shootings to be a response to unrelated violent crime, they need *all* police shootings to be responses to violent crime happening before the police arrived. If this assumption is not valid, all analyses become faulty, because they cannot tease world A (no bias) apart from world B (severe anti-black bias).

The authors' analysis and results are equally consistent with: **“The police only shoot violent criminals, and there is no evidence for bias”**, as they are with: **“police interactions with black people are more likely to result in violence as compared to those with white people”**. Looking at the *actual* data about police shootings (e.g., the Guardian database mentioned by the authors), the second interpretation is much more consistent with reality than the first⁴. If anything, the cases in the Guardian database indicate that the real world is closer to World B than World A.

Although fully reliable data is not available, the Guardian database used by the authors allows us to test whether the assumption the authors' make appears plausible. We coded the most recent three months of the database's data to determine if the initial cause of the stop was related to a violent crime or not. Because there is subjectivity involved in coding and ambiguity in information available, we decided to mark crimes as violent as long as the initial stop was related to a violent crime, no matter how wrong that judgment was or how clearly the shooting was ultimately unjustified. For example, this definition codes the stop of Tyre King, a 13 year old

---

⁴ U.S. Department of Justice data finding black people are more likely to experience force when stopped by police also makes this first interpretation more likely (Davis et al. 2018).

child shot and killed while running away from the police and carrying a BB gun, as a stop starting from a violent crime because officers mistakenly believed King matched the description of an armed robbery suspect. Under this coding scheme, we found that roughly 60% of stops stemmed from violent crimes, and 40% stemmed from non-violent crimes or no crime at all. Our coding manual and descriptions of each case are available at https://osf.io/waqr3/?view_only=21685c4525d349fdbef7cf720cd90b33 if readers wish to evaluate these cases for themselves. However, even if the proportion of non-violent criminals is somewhat different with a revised coding, or varies between month, that does not change the force of our argument. For the authors' method to be valid, *all* stops need to come from violent crime.

One may ask how the authors should have done their analysis⁵, or fault us for not providing a clear alternative answer. That is not our goal. It's clear that there is a huge amount of uncertainty here, and a range of answers are possible. The fact is that the authors cannot answer the question they are asking with the data they have. We agree that this question is important, however we believe the methods used in this paper do more harm than good because this method can lead to conclusions completely conflicting with the underlying data.

---

⁵ While we believe the problem we cover here is the most critical issue with this paper, there are other methodological issues present in this paper including, among others, issues with the reliability of The Counted data (which can use media reports based only on police reports), base rate neglect (for a broader discussion of issues with administrative data on police shootings see Knox et al. 2019), and contradictory results found in other papers which are not adequately explained (Ross 2015). In particular, we would like to note the extensive criticism by Ulrich Schimmack and Rickert Carlsson of a similar publication by many of these same authors (Johnson et al. 2019), which notes a number of serious flaws in that publication as well (Schimmack 2019).

References

Cesario, J., Johnson, D. J., & Terrill, W. (2018). Is There Evidence of Racial Disparity in Police Use of Deadly Force? Analyses of Officer-Involved Fatal Shootings in 2015–2016. *Social Psychological and Personality Science*, 10(5), 586–595.

<https://doi.org/10.1177/1948550618775108>

Davis, E., Whyde, A., & Langton, L. Bureau of Justice Statistics. (2018). *Contacts Between Police and the Public, 2015* (Report NCJ 251145). U.S. Department of Justice.

Desmond, M., Papachristos, A. V., & Kirk, D. S. (2016). Police Violence and Citizen Crime Reporting in the Black Community. *American Sociological Review*, 81(5), 857–876.

<https://doi.org/10.1177/0003122416663494>

Johnson, D. J., Tress, T., Burkel, N., Taylor, C., & Cesario, J. (2019). Officer characteristics and racial disparities in fatal officer-involved shootings. *Proceedings of the National Academy of Sciences*, 116(32), 15877–15882.

<https://doi.org/10.1073/pnas.1903856116>

Knox, D., Lowe, W., & Mummolo, J. (2019). Administrative Records Mask Racially Biased Policing. <http://dx.doi.org/10.2139/ssrn.3336338>

Ross, CT. (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.

<https://doi.org/10.1371/journal.pone.0141854>

Carlsson, R., & Schimmack, U. Comment on Johnson et al. 2019.