

Additional Testimony for the House Committee on the Judiciary
Phillip Atiba Goff, PhD
Inaugural Franklin A. Thomas Professor in Policing Equity
John Jay College of Criminal Justice
and
President and Co-Founder
Center for Policing Equity

Chairman Nadler, Ranking Member Collins, and distinguished Committee Members. It has come to my attention that there is reason for me to submit further testimony to your committee on the subject of race and policing.¹ My goal is to provide the committee with a succinct accounting of what social scientists know about the presence of racial bias in police use of force in the United States.

My task is aided by the fact that we know far less than we should about racial disparities, discrimination, or bias in policing, with more questions than answers. What we do know is that there are large disparities between Black, Native, and Latinx residents on the one hand, and White residents on the other. There is also reason to be confident that racial bias plays a role in producing some portion of those disparities. **To argue that there is a robust science saying there are not racial disparities between White and non-White residents in police use of force is not a serious reading of the literature. To argue that racial biases disadvantage White residents in any robust way is also an unrecognizable distortion of the science.**

In the testimony that follows I will attempt to outline where a serious reading of that literature would suggest a productive conversation might begin. I divide the rest of my testimony into five sections:

- I. Defining Terms**
- II. Levels of Bias**
- III. The State of the Field**
- IV. On “Black-on-Black Crime”**
- V. Closing Thoughts**

I. Defining Terms

For the purposes of this testimony, there are two terms that are critical to define: group-based disparity and bias. A group-based disparity is simply the existence of differences between groups. For instance, people from Philadelphia may watch the television show *The Good Place* at a higher rate than people from Pittsburgh. This does not imply anything about the reason *why* those disparities exist, it simply describes *that* they exist. In the present context, racial disparities in police use of force refer to the greater rate at which Black, Native, and Latinx persons experience police use of force compared to White persons. Calling them disparities does not imply either sinister or virtuous reasons they exist.

¹ I should note that, while I was notified by Keenan Keller that Ms. MacDonald had submitted testimony that made claims about me, as of October 15, neither I nor anyone at CPE received any message from Ms. MacDonald after the testimony I gave on September 19, 2019.

A bias connotes a tendency to favor one group over another. For my purposes in this testimony, bias need not be a product of individual character. It is not synonymous with bigotry. Take for example an employer whose business is on the east side of a town. This employer does not hire people from the west side of town because she fears (rightly or wrongly) that west-siders are more likely to quit. This employer may not dislike west-siders. Still, she worries that the commute from the west to the east side is burdensome making it harder for her business to retain westside residents on staff. While this would not qualify as bigotry in most contexts and for most people, the pattern of hiring that results would still be biased. Crucially, here, biased behaviors are behaviors that are within an individual (or group's) power to alter.

II. Levels of Bias

The research literature on racially disparate police use of force identifies at least three levels at which bias might operate: at the level of cities, at the level of contact rates, and at the level of the decision to use force. For racial bias to operate at the level of cities, it would mean that City X's police department treated people differently than City Y's police department, regardless of whether or not officers *within* Cities X and Y differ in how likely they were to stop or use force on residents of their city. If City X uses force more frequently than City Y and City X has more Black residents than City Y, one might not observe racially biased force *within* either city, but would observe bias *between* cities. While it is reasonable to hypothesize these kinds of city-level differences, vanishingly little research is able to explore them due to a lack of available data.

The second level at which bias could operate is that, within a city, people from Group X are more likely to have contact with police than people from Group Y, but are not necessarily more or less likely for force to be used once police engage them. To understand how bias could work at the level of contact rates, I offer a thought experiment. Imagine that people from Group X are stopped at a rate of 100 stops for every 1,000 group members in a city and people from Group Y are stopped at a rate of 50 stops for every 1,000 group members in a city. If one imagines that police use force in 2% of all contacts with residents, *regardless of any other factors*, then the people in Group X are no more likely to have force used against them than the people in Group Y once they have been contacted by police. But, people in Group X are still *twice as likely* to have force used on them overall. In other words, the rates of contact lead to racial disparities, but not officer decisions once contact is made. Sadly, this too is exceedingly difficult to study. Without good proxies for police contact, researchers are mostly left to use unsatisfactory substitutes or highly sophisticated quantitative methods to impute estimates of police/community contacts.

Finally, there is the potential for bias at the level of the decision to use force. In popular media, it is common for this to be discussed as if it is the only point at which bias might influence outcomes. If officers under similar conditions show no evidence of bias, the reasoning goes, then there can be no bias in policing. This is wrong on its face and it is crucial for members of the committee—and the public—to understand why. Even if officers do not demonstrate bias under similar conditions, bias can still influence police decisions to use force at the rate of contact level or the city level. That is, patrol deployment policies and/or department-wide

policies and cultures can be biased and produce racially disparate outcome independent of officer decisions and would be no less damaging to the communities they afflict.

In fact, if non-White residents are disproportionately targeted for police contact, it is reasonable to hypothesize that there would be a lower rate of police force towards them because they are contacted for less serious reasons. Again, this could lead to higher rates of force against non-White residents, just at the level of contact, not officer decision.

The scientific literature has focused disproportionately on bias at the level of officer decision because the data that are available are most robust here. The disproportionate focus on officer-level decisions is also part of the reason there has been such serious skepticism around bold claims ruling out bias in officer force. Any research paper that attempts to rule out bias without considering at least two of the three levels outlined above is an overreach of the data. This is the error to which Johnson, Cesario, and colleagues admitted in their response to public criticism, particularly that of Knox and Mummolo. From their response (already entered into the congressional record) in which they respond to the criticism that their paper misrepresents their data as being able to offer insight into the rates at which Black and White residents are shot:

Part of the reason for this misconception is likely due to the language used in our significance statement. We wrote ‘White officers are not more likely to shoot minority civilians than non-White officers.’ We should have written this sentence more carefully as it currently refers to estimating $\Pr(\text{shot}|\text{race}, X)$. What we should have written was a sentence about what we did estimate— $\Pr(\text{race}|\text{shot}, X)$: As the proportion of White officers in a FOIS increased, a person fatally shot was not more likely to be of a racial minority. This was our mistake, and we appreciate the feedback on this point. (Johnson & Cesario, 2019, pp. 1)

The statement “White officers are not more likely to shoot minority civilians than non-White officers” is not supported by the data made available in the paper despite this being the central claim in the paper’s significance statement. This is also what I meant during my testimony before this committee on September 19, 2019.

This critique, overstating findings about disparities in police use of force to draw conclusions about the existence of bias at all levels of policing, is also the central critique of Fryer’s research, which was also referenced during testimony before this committee on September 19, 2019. While Fryer’s work on policing has been vociferously critiqued both in public and among scholars, those critiques are perhaps best encapsulated by the work of Ross and colleagues. In their recent paper (Ross, Winterhalder, & McElreath, 2018), Ross and colleagues are able to resolve the seemingly aberrant findings of Fryer with the vast majority of the literature by demonstrating how Fryer’s findings may be an example of Simpson’s paradox: the tendency to see the relationship between two variables in a small sample reverse in the larger population. Again, because Fryer was not able to (and neither attempted to nor claimed to) test for bias at the level of contact or cities, his claims about the absence of anti-Black bias (or even the presence of anti-White bias) were met with considerable scientific skepticism.

My understanding from conversations with his colleagues at the time is that the public backlash against this work troubled Fryer, and he emphasized in public appearances that the deadly force findings were only part of his analyses. While he was still permitted on the campus, his colleagues characterized his reaction to the controversy surrounding this paper as being embarrassment and frustration that people were “using his research” to argue there is no anti-Black bias in police violence. His dubious conclusions around deadly force aside, Fryer’s paper is consistent with the vast majority of previous research in demonstrating that racial disparities persist, even controlling for crime and poverty. In other words, neither crime nor poverty are sufficient to explain racial disparities that harm Black residents in police use of force (see Goff et al., 2016). To say that Fryer found evidence of anti-White bias in police use of force is—at best—a selective reading of his full research report. Because this claim has been roundly critiqued by scholars and in public means that anyone claiming, without caveat, that Fryer’s work is conclusive evidence of anti-White bias in policing has either not been exposed to the literature or is willfully ignorant of it.

Similarly, Correll and colleagues (also referenced during the hearing on September 19, 2019) report robust anti-Black bias in their own analyses, with Correll documenting those biases across multiple studies (Correll et al., 2007; 2011; Sadler, Correll, Park, & Judd, 2012). Specifically, officers’ speed of response is consistent with negative stereotypes about non-Whites and positive stereotypes about Whites across all reported studies, even when officer “errors” in shooting are not.

Because I have had reason to focus on the work of Cesario and Fryer, it may seem from reading this testimony that there is sizable debate in the field about the role and direction of bias in police use of force. I believe that impression would be a distortion of the literature. While these two papers highlight how methodological or conceptual shortcomings of a given piece of research can be misleading, they are exceptions to the norm in this research literature. In fact, there is a great deal of agreement in the literature, which is where I turn my attentions next.

III. The State of the Field

After mentioning to my colleague Jonathan Mummolo that I would be preparing additional congressional testimony, he shared a project that he and a student, Tori Gorton, recently completed. This project searched for articles containing the keywords “racial bias,” “police,” “force” and “data” in order to identify the general tenor of the research literature on racial bias in police use of force. Of the roughly 11,000 articles their search returned, they reviewed the 50 articles with the most citations. Of those 50 articles, 92% (46) reported some form of bias against non-White populations. A total of seven studies reported some form of anti-White bias, but only one of those 50 reported anti-White bias without *also reporting bias against non-White populations*.

Those numbers do not represent a robust debate.

In other words, **research claiming to reveal robust patterns of anti-White bias is exceedingly rare in this literature and often complicated by an inability to look at multiple levels at which bias might operate**. As I stated at the beginning of my testimony, there is far too

little that social scientists know about racial bias in policing. I was part of the National Academies of Sciences committee that examined racial disparities resulting from pro-active policing, which came to the consensus conclusion that we know far less than we should (National Academies of Sciences, Engineering, and Medicine, 2018). But we do know that there are robust racial disparities in police behaviors that disadvantage non-White communities and there is good reason to suspect that bias plays at least some role in the production of those disparities. Mummolo and Gorton's work reflects my own sense of the field: that there are few research articles that find evidence of anti-White bias, that the ones that do are methodologically challenged, and that they receive outsized attention when compared to their relative infrequency in the literature.

IV. On "Black-on-Black Crime"

Because it comes up so frequently, I feel compelled to include a short note on the notion of "Black-on-Black crime" being a reasonable explanation for racial disparities in police outcomes. Importantly, there are two ways to read references to "Black-on-Black crime" in debates about racial disparities in policing. The first is a bad faith attempt to cast blame on Black communities for negative outcomes they experience. The second is a serious methodological concern. Perhaps police are simply going to where crime is and that drives racial disparities in police outcomes? This can feel to some like a distasteful question to ask given how frequently some feel the term "Black-on-Black crime" is used in bad faith. Still, it would be unreasonable, scientifically, to ignore the possibility that crime rates drive patrol deployments. And, if rates of contact predict rates at which police use force, crime rates likely play some role in the responsible analysis of racial disparities in police outcomes.

The good news is, researchers have included this concern in their analyses. Scientists have long taken seriously the idea that police are genuinely driven in part (or even entirely) by a desire to respond to serious crime. Consequently, the best and most cited research articles on racial disparities in police outcomes account for police contact by factoring in arrests, crime, stops, citations, or some other proxy for contact and/or illegal activity.

And the best attempts to account for crime and poverty, again, reveal that they are not sufficient to explain racial disparities in police outcomes (Gelman, Fagan, & Kiss, 2007), including force (Goff et al., 2016). In other words, "Black-on-Black crime" is simply not sufficient to explain the rates of disparities in police stops or use of force. My experience working with law enforcement tells me that explaining these findings is essential to productive engagements between police and the communities they are sworn to protect. At the very least, after it is explained, it is clear that no one could reasonably appeal to "Black-on-Black crime" in good faith. This, itself, is often progress.

V. Closing Thoughts

As I said at the beginning of this testimony, there are far more questions about racial bias in policing than there are answers. What psychological vulnerabilities best predict biased behaviors? What is the role of bigotry in racial disparities? What situations best protect officers against the influence of bias on their behaviors—and therefore communities against those biases?

Scientists are just beginning to learn the answers to these questions and hundreds more due to the previous lack of available data and substandard methodologies. But there are some questions that those who are competent to evaluate the science can agree we have answered. And it is not productive to pretend otherwise nor to allow those who are not competent to evaluate the science to position themselves as if they are.

We know that there are racial disparities in police use of force that are detrimental to Black, Native, and Latinx communities. We are reasonably confident that bias accounts for some portion of those disparities (though, again, this is not a statement about police character, *per se*). As a result, we should allow serious conversations about what to do with these problems to begin from there. To do otherwise does no favors to those who wear the badge or the communities that most need them to be excellent at their jobs.

References

- Correll, J., Park, B., Judd, C. M., Wittenbrink, B., Sadler, M. S., & Keesee, T. (2007). Across the thin blue line: police officers and racial bias in the decision to shoot. *Journal of personality and social psychology*, *92*(6), 1006.
- Correll, J., Wittenbrink, B., Park, B., Judd, C. M., & Goyle, A. (2011). Dangerous enough: Moderating racial bias with contextual threat cues. *Journal of experimental social psychology*, *47*(1), 184-189.
- Gelman, A., Fagan, J., & Kiss, A. (2007). An analysis of the New York City police department's "stop-and-frisk" policy in the context of claims of racial bias. *Journal of the American Statistical Association*, *102*(479), 813-823.
- 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.
- Ross, C. T., Winterhalder, B., & McElreath, R. (2018). Resolution of apparent paradoxes in the race-specific frequency of use-of-force by police. *Palgrave Communications*, *4*(1), 61.
- Sadler, M. S., Correll, J., Park, B., & Judd, C. M. (2012). The world is not black and white: Racial bias in the decision to shoot in a multiethnic context. *Journal of Social Issues*, *68*(2), 286-313.