

Reconciling Results on Racial Differences in Police Shootings

By ROLAND G. FRYER, JR.*

Police use of force – particularly lethal force – is one of the most divisive issues of the twenty-first century. When an officer discharges a weapon, it’s a life changing event. A single bullet – which weighs about .02 pounds and is 10mm long – can end a life, erase a pension, or change the image of those who are sworn to serve and protect.

To understand the nexus of race, criminal justice, and police brutality, academics and journalists have begun to amass impressive datasets on Officer-Involved-Shootings (OIS). The Washington Post (2015) provided important stylized facts about individuals killed at the hands of police officers in one calendar year. Police fatally shoot nearly 1,000 individuals a year in the U.S. As the report argues, the incidents that have ignited protests – white police officers killing unarmed black men – represent less than four percent of fatal police shootings.

The Guardian (2015) collected similar data in their much publicized series, *The Counted*, and tallied 1,134 deaths of Americans at the hands of police in 2015. And, similar to The Post’s database, The Guardian emphasizes that although black males aged 15-34 make up 2 percent of the population they comprise 15 percent of all deaths logged in 2015 – five times the comparison number for whites. One in every sixty-five – or 1.5 percent – deaths of a black man in the US is a killing by police. Roughly twenty-five percent of blacks that were killed were unarmed, eight percentage points higher than whites (*Guardian*, 2015).

More recently, VICE News collected data on OIS from 47 of the 50 largest police departments in the U.S (*VICE*, 2017). The major difference between the VICE

data and that assembled by The Post or Guardian is that it counts all OIS – fatal or non-fatal – whereas The Post and Guardian data count fatal shootings only. VICE catalogs 4,400 individuals involved in OIS between 2010 and 2016. Two-thirds of the individuals shot by police survived. Twenty percent of people involved in OIS were unarmed – four times as large as the rate of unarmed people documented in The Post/Guardian data of fatal shootings. Moreover, blacks were 55 percent of shooting victims, which is more than double the share of the black population in these jurisdictions. Black subjects shot by police were more likely to be shot during encounters that began with a traffic or pedestrian stop. They were no more likely to be armed with a gun than whites but less likely to be armed with any weapon (*VICE*, 2017).

Among social scientists, the literature is surprisingly thin. Ross (2015) estimates racial differences in county-specific relative risks of being shot by police, using crowd-sourced data from the US Police-Shooting Database. He estimates that the probability of being black, unarmed, and shot by police is nearly 3.5 times larger than the probability of being white, unarmed, and shot by police.

In stark contrast, Fryer (*forthcoming*) finds that, conditional on a police interaction, there are no racial differences in OIS on either the extensive or intensive margins.¹ Using data from Houston, Texas – where I have both OIS and a randomly chosen set of interactions with police where lethal force may have been justified but was not used – I find, after controlling for suspect demographics, officer demographics, encounter characteristics, suspect weapon and year fixed effects, that blacks are 27.4

* Fryer: Harvard University and NBER, 1805 Cambridge Street, Cambridge MA 02138, rfrayer@fas.harvard.edu. I would like to thank Tanaya Devi, Hannah Ruebeck, and Damini Sharma for exceptional research assistance. All errors are my own.

¹Goff, et al. (2016) reports similar findings on the application of lethal force, but their data is not available for public use and thus cannot be included in this review.

percent *less* likely to be shot at by police relative to non-black, non-Hispanics. Investigating the intensive margin – who shoots first in an encounter with police or how many bullets were discharged in the endeavor – there are no detectable racial differences.

The key takeaway suggested by The Post (2015), The Guardian (2015), VICE (2017) and Ross (2015) is that there are large racial differences in OIS that are evidence of racial bias or, at the very least, problematic. As the Post writes, “Although black men make up only 6 percent of the U.S. population, they account for 40 percent of the unarmed men shot to death by police this year.” The Guardian begins their article remarking, “Young black men were nine times more likely than other Americans to be killed by police officers in 2015,” and VICE exclaims, “Black people were shot more often and at higher rates than people of any other races.” Ross (2015) concludes, “There is evidence of significant bias in the killing of unarmed black Americans relative to unarmed white Americans.”

Yet, these descriptive statistics, while poignant, do not prove racial bias. Indeed, when I compare the different data sets described above using a common estimating equation I conclude that the “evidence for bias” in some studies and not others is mainly a result of misspecified regression equations, not inherent differences in the datasets.

I. The Simple Econometrics of Estimating Racial Differences in Officer-Involved-Shootings

Conceptually, estimating racial differences in OIS is very similar to the large and well-established literature in labor economics designed to understand whether racial differences in wages is due to discrimination.

Borrowing extensively from that literature and simplifying to the extreme, let’s suppose there are only two races: black and white. Let $R_i = 1$ if individual i is black and $R_i = 0$ if they are white. And, suppose that Y_i^1 denotes the value of Y (e.g.

probability of being involved in an OIS) if individual i is black and Y_i^0 denote the value of Y if individual i is white. Then, the “race effect” for individual i is $\tau = Y_i^1 - Y_i^0$ – that is, the difference in Y that can be attributed to an individual’s race. This quantity is the proverbial “holy grail” – the parameter that we are all attempting to estimate but never quite do.

Consider the following linear regression equation: $Y_i = \beta_0 + \tau R_i + \beta_1 X_1 + \dots + \beta_n X_n + \epsilon_i$. One can trivially demonstrate that this regression recovers the “race effect” if any difference between the groups that we’d observe in the absence of bias is zero. To see this formally, write Y_i as $Y_i^0 + R_i(Y_i^1 - Y_i^0)$. Note: τ measures $E[Y_i | R_i = 1] - E[Y_i | R_i = 0]$. Using these two facts:

$$\tau = \underbrace{E[Y_i^1 - Y_i^0]}_{\text{Race Effect}} + \underbrace{E[Y_i^0 | R_i = 1] - E[Y_i^0 | R_i = 0]}_{\text{Selection Effect}}.$$

It is straightforward to see that estimating Equation (1) recovers the race effect when the selection effect is zero. This happens when $Y_i^1, Y_i^0 \perp R_i | X_i$, which ensures $E[Y_i^0 | R_i = 1] = E[Y_i^0 | R_i = 0]$. In words, a linear regression can recover the “race effect” if race is “as good as randomly assigned,” conditional on the covariates (X ’s). When the selection effect is not zero then one can misinterpret the mis-measurement of X ’s (e.g. omitted contextual variables) as racial bias. One way to ensure that race is “as good as randomly assigned” is to implement an experiment (see Bertrand and Mullainathan (2004) for a classic example). None of the studies mentioned above meet this threshold.

II. Analyzing the Literature

In this section, using the above conceptual framework, I analyze the growing literature on racial differences in OIS on a set of five rubrics: (1) Reliability of Police Data; (2) Representativeness of Data; (3) Endogeneity of Police-Civilian Contact; (4) Relevant Counterfactuals; and (5) Quality of Covariates.

The brilliance of the datasets constructed by the staffs at The Post and Guardian is that they do not rely heavily on descriptions by police. Their goal is to provide a more accurate number of American citizens whose lives were ended as a result of a police shooting. Because of their exhaustive search process – which included news reports, public records, and original reporting – the data seem representative for what it is: a count of the fatal shootings by police in America in a given year. The Post and Guardian documented more than twice as many fatal shootings in 2015 as the annual average annual reported by the FBI over the past decade.

These data do not deal with the endogeneity of contact by police, but assuming blacks are more likely to be unfairly targeted than whites, the racial differences exhibited in their data would only be larger if one did account for the fact that police-civilian interactions are themselves (potentially) biased.

However, two of the most important rubrics for internal validity are not consistent with The Post and Guardian datasets – having a relevant counterfactual and the quality of covariates. Theoretically, The Post’s data is rich enough to account for several different potential counterfactuals. For instance, one could estimate the probability of being killed if unarmed and black relative to the probability of being killed if unarmed and white, after accounting for the context of the fatal encounter. Yet, the more relevant counterfactual – given it is less likely to be plagued by selection into the sample – is to estimate when police decide to pull the trigger. For this, one needs to catalogue situations in which police would have been legally justified in using lethal force but chose not to – along with the race of the suspect and other contextual factors.

Both The Post and Guardian collect manner of death, whether or not the suspect was armed, weapon type, demographics, and the city in which the shooting took place. The Post additionally collects whether there were signs of mental illness, the threat level to officers, and whether the officer was wearing a body camera. The

Guardian additionally records the address at which the shooting took place and the name of the law enforcement agency responsible.

To draw firm conclusions about the “race effect” one needs to assume that once this sparse set of variables is accounted for, race is “as good as randomly distributed” and that there are no other differences in suspect behavior or any other potential contextual factor that is important. With tremendous respect to the herculean efforts of The Post and Guardian for collecting new data on one of the most important social issues of our time, this assumption is ludicrous and one can make no such claims of any “race effect” using this data.²

A typical retort to the lack of rigor when employing descriptive statistics is to observe the size of the racial difference and posture that such large disparities are not likely to be explained by differences in context. That type of argument is more about clinging to deeply held beliefs than clinging to a valid research design. There is no way *a priori* to understand how much of the variance in an outcome might be absorbed by covariates without first analyzing the data.

The VICE data has similar advantages and disadvantages, but is more comprehensive in that it includes both fatal and non-fatal shootings from 47 of the 50 largest police departments in America – which is in itself phenomenal – but less representative than The Post or Guardian datasets. Their set of potential covariates, which includes whether or not the suspect was armed, suspect and officer demographics, the nature of the stop, and number of officers, is quite limited and they do not have any data on situations in which lethal force would have been justified but was not used.

Recall, Ross (2015) estimates racial differences in county-specific relative risks of being shot by police, controlling for whether or not the civilian was armed, and concludes that the results provide evidence of significant bias. An issue

²Nix et al. (2017), using the data in *Post* (2015), employ this approach – accounting for a small set of covariates and claiming the residual is “implicit bias.”

with the crowdsourced U.S. Police-Shooting Database used in Ross (2015) is that it omits large fractions of shootings that occurred and it's not clear which types are recorded. To illustrate this, there are 10 locations in Ross (2015) that are also in Fryer (*forthcoming*). Of the 606 officer-involved shootings in these locations, between 2011 and 2014, noted in Fryer (*forthcoming*), only 102 of them (17 percent) were reported in Ross (2015). He cannot account for the endogeneity of police-civilian contact, but again, this is likely to only increase estimates of racial differences in the county-specific relative risk. A strength of the Ross analysis is that the crowdsourced data may be less likely to be biased by police department reporting practices.

To assemble the database utilized in Fryer (*forthcoming*), I requested police narratives for all OIS from 10 large police departments. Houston Police Department also provided a random sample of narratives for events that did not end in an OIS but in which justifiable force could have been used. Narratives were twice-coded by research assistants for relevant variables.

This database addresses many of the issues that plague the literature – there is a relatively comprehensive set of contextual variables (290, for instance, hand-coded from Houston arrest data), and the counterfactuals are less likely to be biased by selection into sample – but the data are not nearly as representative (only 10 police departments) as data collected by The Post, The Guardian, or VICE. Indeed, the analysis in Fryer (*forthcoming*) is tantamount to trying to estimate discrimination in a labor market by analyzing data from firms willing to supply it to researchers. Put differently, while the estimates in Fryer (*forthcoming*) are more internally consistent, their external validity is unknown.

There are two important fault lines in Fryer (*forthcoming*): (1) the endogeneity of police-civilian contact and (2) reliability of police department reports. To the extent that there are racial differences in the probability of an interaction with police, the data in Fryer (*forthcoming*) may

omit a very important margin.³ One may discover no differences in use of lethal force, *conditional on an interaction*, even if there are large racial differences in the probability of the types of interactions in which lethal force may be used.

Understanding racial differences in the probability of police interaction is difficult. One has to account for differential exposure to police, race-specific crime participation rates and perhaps most importantly, pre-interaction behavior that civilians exhibit. Ideally, one might set up a field experiment that randomly assigns similar individuals (across all physical dimensions except race) to the vicinity of the same patrolling officers in a neighborhood and instruct them to behave identically. Conditional on random assignment, identical behavior, and race-specific crime rates, any differences in the probability of interaction could be interpreted as racial bias in police stopping behavior. Without ideal data, researchers often compare the racial distribution of stopped civilians to the racial distribution of various “at risk” civilians that could potentially be stopped. Determining the probability of an interaction is essentially a search for the correct “risk set.” The evidence in Fryer (*forthcoming*) offers the following rough rule of thumb – if one assumes that police are non-strategic in their stopping behavior, then there is clear bias. Conversely, if one assumes that police are stopping individuals they are worried will engage in violent crimes, the evidence for bias is small.

A second concern is the reliability of police department reports. There are two types of potential bias. First, police officers may bend the truth about the context of a particular interaction so as to justify their own actions; for instance, indicating a suspect was threatening when they were calmly following an officer's commands. This type of bias is less of a concern in Fryer (*forthcoming*) because the qualitative results are identical whether or not one includes contextual factors about the

³This is also true of all other datasets collected, but is more of an issue when one does not find bias.

encounter recounted by police. A second type of bias is that officers may be more likely to charge black suspects with crimes such as resisting arrest or attempted assault on a public safety officer rather than misdemeanors, relative to whites, for identical behavior. This type of bias is an important limitation of Fryer (*forthcoming*) because it implies that the counterfactuals coded from arrest data may themselves contain bias. It is unclear how to estimate the extent of such bias or how to address it statistically.

III. Reconciling the Literature

Datasets created by The Post, Guardian, VICE, or the US Police-Shooting Database provide important descriptive statistics about OIS in America, but have severe limitations for determining whether those statistics are due to racial bias. Fryer (*forthcoming*) is better suited for this type of analysis, but whether and how much one should extrapolate from 10 cities willing to supply data is unknown.

I replicate some of the key findings in each of the publications discussed throughout with the data from the alternative data. That is, I recreate the descriptive statistics used by The Post and Guardian, VICE, and Ross (2015) using the data from each publication and Fryer (*forthcoming*). This allows one to better understand whether differences across datasets are due to different statistical decisions or inherent differences in data. For instance, if the data used in Fryer (*forthcoming*) – the most detailed but least representative – shows less bias when one estimates the fraction black unarmed who are killed by police relative to the fraction white unarmed, then this further calls into question the generality of the data.

Table 1 presents these estimates. The descriptive statistics presented by The Post (2015), The Guardian (2015), VICE (2017), and Ross (2015) are remarkably consistent across all datasets – including Fryer (*forthcoming*). For instance, The Guardian (2015) shows that the likelihood of black men between the ages of 15-34 to be killed by police, relative to all other races, is 9.82.

Using the data collected by VICE this number is 9.82, 15.93 when using the data in Ross (2015), and 14.87 when using the data in Fryer (*forthcoming*).

Perhaps the most striking finding is when one replicates the analysis in Ross (2015) across all five datasets: calculating the probability of being black, unarmed, and shot by police divided by the probability of being white, unarmed, and shot by police. A quantity greater than one is consistent with racial bias. Using the data from Ross (2015), this ratio is 3.28. Using the data from the Post database I get 6.20 and 5.99 if using the data in Fryer (*forthcoming*). In other words, if I ignore the detail available in the Fryer data and simply report the descriptive statistics reported in Ross (2015), I could conclude that the data provided evidence of even more racial bias than that reported in Ross (2015). Yet, when using the simple statistical framework that economists have used for more than a half century to analyze racial differences on myriad dimensions – from wages to incarceration to teen pregnancy – the evidence for bias disappears. The differences in results on police shootings in America seem to be driven by differences in what qualifies for a valid research design and not differences in datasets.

IV. Conclusion

The time has come for a national reckoning on race and policing in America. But, the issues are thorny and the conclusions one can draw about racial bias are fraught with difficulty. The most granular data suggest that there is no bias in police shootings (Fryer (*forthcoming*)), but these data are far from a representative sample of police departments and do not contain any experimental variation. We cannot rest. We need more and better data. With the advances in natural language processing and the increased willingness of police departments to share sensitive data, we can make progress.

For those of us who desire a more perfect union, police use of force has become our Gettysburg. Of course, black lives matter as much as any other lives. Yet, we

TABLE 1—RECONCILING RESULTS ON RACIAL DIFFERENCES IN POLICE SHOOTINGS

	Post (1)	Guardian (2)	Vice (3)	Ross (4)	Fryer (5)
Paper: Statistic					
<i>Washington Post</i> : Percent of black civilians among unarmed men killed by police	39.56	34.25	52.82	38.71	43.84
<i>Guardian</i> : Likelihood of black men between the ages of 15-34 to be killed by police compared to others	9.67	9.82	9.82	15.98	14.87
<i>Vice</i> : Percent unarmed in sample	9.45	20.42	21.31	28.29	20.50
Percent black civilians in sample	26.80	27.21	54.59	34.95	45.70
<i>Ross (2015)</i> : Prob. of being black, unarmed and shot by police relative to prob. of being white, unarmed, and shot by police	6.20	3.44	3.21	3.28	5.99

In row (1), I calculate the likelihood as the ((Fraction of 15-34 black men killed by police/Fraction of 15-34 black men in population)/(Fraction of everyone else killed by police/Fraction of everyone else in population)). In row (2), I calculate the percentage of black civilians among unarmed men killed by police. In row (3), I calculate the percentage of unarmed civilians and black civilians in the sample. In row (4), I use the companion program to Ross (2015), available in *R*, to calculate the county-level risk ratios. To calculate this statistic, I collapse each OIS-level dataset to the county level (Post and Guardian) or the police department level (VICE and Fryer (*forthcoming*)). To collapse cities to the county level (Post and Guardian), I assign cities to the county that they fall the majority in. The VICE and Fryer (*forthcoming*) datasets provide the police department responsible for the OIS. I collapse these counts by police departments and assign each police department the population of the city or county that it has jurisdiction over. I take city and county populations by race from the 2011 ACS 5-year estimates. The Ross (2015) estimates may contain simulation error but are replicated using the exact program used in Ross (2015). In addition, the exact results are not reproducible because the program does not set a seed.

do this principle a disservice if we do not adhere to strict standards of evidence and take at face value descriptive statistics that are consistent with our preconceived ideas. ‘Stay Woke’ – but critically so.

REFERENCES

- Arthur, Rob, Taylor Dolven, Keegan Hamilton, Allison McCann, and Carter Sherman. 2017. “Shot by cops and forgotten.” *VICE News*, December 11. news.vice.com/en_us/article/xwvv3a/shot-by-cops
- Bertrand, Marianne and Sendhil Mullainathan. 2004. “Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination.” *American Economic Review* 94 (4): 991-1013.
- Fryer, Roland G. *Forthcoming*. “An Empirical Analysis of Racial Differences in Police Use of Force.” *Journal of Political Economy*.
- Goff, Philip A., Tracey Lloyd, Amanda Geller, Steven Raphael, and Jack Glaser. 2016. “The Science of Justice: Race, Arrests, and Police Use of Force.” *Center for Policing Equity*.
- Kindy, Kimberly, Marc Fisher, Julie Tate, and Jennifer Jenkins. 2015. “A Year of Reckoning: Police Fatally Shoot Nearly 1,000.” *The Washington Post*, December 26. www.washingtonpost.com/sf/investigative/2015/12/26/a-year-of-reckoning-police-fatally-shoot-nearly-1000/
- Nix, Justin, Bradley Campbell, Edward Byers, and Geoffrey Alpert. 2017. “A Bird’s Eye View of Civilians Killed by Police in 2015.” *Criminology & Public Policy* 16 (1): 309-340.
- Ross, Cody. 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).
- Swaine, Jon, Oliver Laughland, Jamiles Larteyand, and Ciara McCarthy. 2015. “Young Black Men killed by US Police at Highest Rate in Year of 1,134 Deaths.” *The Guardian*, December 31. www.theguardian.com/us-news/2015/dec/31/the-counted-police-killings-2015-young-black-men