

Lethal Force in Black and White: Assessing Racial Disparities in the Circumstances of Police Killings

Shea Streeter, Stanford University

African Americans are nearly three times as likely to be killed by police as whites. This paper examines whether this racial disparity is due in part to racial differences in the circumstances of police killings. To assess whether and how these circumstances predict the race of a decedent, I use machine learning techniques and a novel data set of police killings containing over 120 descriptors. I find that decedent characteristics, criminal activity, threat levels, police actions, and the setting of the lethal interaction are not predictive of race, indicating that the police—given contact—are killing blacks and whites under largely similar circumstances. The findings suggest that the racial disparity in the rate of lethal force is most likely driven by higher rates of police contact among African Americans rather than racial differences in the circumstances of the interaction and officer bias in the application of lethal force.

After his son was killed by police in a shopping mall, an Alabama father told reporters: “If he had been white they wouldn’t have shot him” (Shah 2018). Similar assertions have been made following a number of high-profile police killings in recent years, particularly in situations where the individual was black and killed under seemingly innocuous circumstances (Chan 2016; *Guardian* 2017; Swaine 2014). In some of these now infamous cases, decedents have allegedly been unarmed, shot in the back, had their hands up, or killed with their children present.¹ Furthermore, new research has found that blacks are nearly three times as likely as whites to be killed by law enforcement, reinforcing a perception that the police killings of black Americans are uniquely unjust (Buehler 2017; DeGue, Fowler, and Calkins 2016). Scholars are still divided, however, over the causes of this threefold racial disparity, and until recently there were not adequate data to determine whether the police kill whites and blacks under different circumstances. In order to further investigate these popular perceptions and empirical trends, this paper examines whether, and how, the circumstances of police killings vary by race.

To assess whether there are racial differences in the circumstances of lethal force incidents, I employ machine learning techniques on an original data set of nearly 1,200 nationwide police killings in 2015. Using a sample of this new data, I analyze over 120 hand-collected variables related to the actions and attributes of the decedent, the police response, and the setting of the lethal interaction. The large number of potential predictors and interactions preclude the use of conventional methods like linear regression; however, machine learning techniques provide a solution for investigating the racial patterns in this high-dimensional data. In the analysis I use a variety of methods to test whether it is possible to predict the race of decedents based solely on the circumstances of the police killing. These methods also allow me to identify which set of predictors most differentiate the fatalities of white Americans from those of African Americans.

The results from the machine learning algorithms reveal that, despite large disparities in the rate of lethal force, there are not systematic racial differences in the observable circumstances of these lethal encounters. In other words, factors like being unarmed, threatening the police, or being engaged in

Shea Streeter (sheas@stanford.edu) is a PhD candidate at Stanford University, Stanford, CA, 94305.

This article was solicited by the editors, was presented and critiqued in a workshop, and received one external review. Data and supporting materials necessary to reproduce the numerical results in the article are available in the *JOP* Dataverse (<https://dataverse.harvard.edu/dataverse/jop>). An online appendix with supplementary material is available at <https://dx.doi.org/10.1086/703541>.

1. These circumstances were present in high-profile lethal incidents involving Eric Garner (2014), Walter Scott (2015), Michael Brown (2014), and Philando Castille (2016), respectively.

The Journal of Politics, volume 81, number 3. Published online June 7, 2019. <http://dx.doi.org/10.1086/703541>
© 2019 by the Southern Political Science Association. All rights reserved. 0022-3816/2019/8103-0029\$10.00

criminal behavior do not reveal substantive information about whether the decedent is likely to be black or white. Even with a large number of predictors and thousands of model specifications, advanced statistical methods grant only marginally more predictive power than naively assigning race to decedents. This indicates that, conditional on contact, the police are killing whites and blacks under largely similar circumstances. Bearing in mind the controversial killings of African Americans, these findings suggest that the police are also killing a similar proportion of whites under dubious circumstances, though such cases have not garnered equivalent notoriety.

The findings of this paper also provide insight into the possible source of the racial disparity in the rate of lethal force. In the next sections, I explain how analyzing the circumstances of police killings can aid in determining whether the disparity originates from differences in police contact, differences during the interaction, or from officer bias in the application of lethal force. Given the empirical findings, I argue that the most probable explanation for the disparity is that African Americans are coming in contact with police at higher rates than whites, leading to a much greater risk of being killed from police encounters. However, I also find that the types of police contact that social scientists most often study—discretionary contact like traffic stops and street stops—only account for a small proportion of lethal incidents. Consequently, I conclude the paper by discussing how the results point to new avenues for research into lethal force incidents as well as police-civilian contact more broadly.

CONTRASTING EXPLANATIONS: RACIAL BIAS OR DISPARATE CONTACT

Scholars across the social sciences have been working to identify why the relative risk of being killed by police is so much higher for African Americans than for whites. As a result of this work, there are two competing explanations for the source of this racial disparity. The first theory is that implicit racial biases influence officers' decisions to shoot, leading them to kill a disproportionate number of black people. In contrast, some recent work has countered that there are no racial differences in the use of lethal force but that the disparity in the rate of police killings is driven by higher minority contact with the police.

The research on implicit racial bias and its applications to policing has come primarily from the field of social psychology. A number of experimental studies have found that in the dominant culture, blackness has come to be associated with criminality and threat. For example, one study found strong, unconscious associations between black faces and crime-related words and images (Eberhardt et al. 2004). In the same publication, a second experiment revealed that police officers

consistently ranked faces with more stereotypically black features as more "criminal looking." Troublingly, this anti-black bias is so strong that the linkage between blackness, threat, and criminality extends even to black children as young as five years old (Todd, Thiem, and Neel 2016).

Studies have also found that this perception of blackness as threatening can have direct effects on the decision to shoot criminal suspects. In video simulations, both civilians and police officers were quicker to shoot armed suspects who were black and were more likely to mistakenly shoot unarmed blacks than unarmed whites (Correll et al. 2002, 2007).² A conclusion from this literature is that officers' misperception and potential overreaction to the threat that African Americans pose could be driving racial disparities in the rate of lethal force.

A second line of research indicates that the disparity in the rate of police killings is due to racial differences in the frequency of police contact. A number of studies have demonstrated that blacks are more likely to come in contact with the police due to traffic stops, street stops, and even from SWAT raids (Baumgartner et al. 2017; Eckhouse 2018; Gelman, Fagan, and Kiss 2007; Lerman and Weaver 2014; Lundman and Kaufman 2003; Mummolo 2018). The probability of contact is also higher because officers are often more heavily deployed in minority neighborhoods (Holmes et al. 2008; Meehan and Ponder 2002; Stults and Baumer 2007). Although these studies are usually quite geographically restricted, the consistent findings across regions and forms of contact indicate that there are likely concrete racial differences in the frequency of police interactions.

Leveraging new data on lethal and nonlethal police contact, recent studies have argued that there are actually no racial differences in officers' decisions to shoot civilians. Conditional on a person being stopped by police, researchers have found that blacks are no more likely to be killed by law enforcement than whites (Fryer 2019; Miller et al. 2017; Worrall et al. 2018). In other words, these studies challenge the earlier research on implicit racial bias and suggest that the experimental findings may not translate into observable racial differences in the real world application of force. Miller et al. (2017) explain that while the probability of being stopped by police is higher for minorities, the relative risk that they will be killed or injured during those interactions is equivalent across racial groups.

2. This bias was slightly reduced for police officers, which the researchers attributed to their extensive training compared to ordinary civilians.

DEFINING THE LETHAL INTERACTION PROCESS

An important shortcoming of the existing literature is that it only examines two stages of what I term the “lethal interaction process”: the series of actions and circumstances that result in a lethal force incident. The process begins with a police encounter in what I term the “contact stage” and ends with an officer’s decision to use lethal force in the “trigger stage.” To oversimplify the conclusions of the literature: either implicit racial bias prompts officers to shoot African Americans more frequently in the trigger stage, or police officers shoot without bias but encounter blacks more frequently in the contact stage.³ The primary theoretical innovation of this paper is that I introduce a third, intermediary stage into the interaction process which can help us estimate the likely source of the racial disparity.

The “interaction stage” encompasses everything that happens during the police-civilian interaction after the contact stage and leading up to the use of lethal force. An almost infinite number of events and circumstances could occur during this stage. Officers have many tactics at their disposal and could issue verbal commands, use nonlethal tactics, go undercover, or deploy a SWAT team. The civilian in this interaction could also engage in any number of activities including compliance, criminal activity, or pointing a weapon at officers. On a macro level, this interaction could occur during the day or night, in front of witnesses or in an abandoned alley, as the result of a traffic stop or a service call.

The set of circumstances and events that occur during the interaction stage should influence the probability that an officer will use lethal force during the trigger stage (Bolger 2015). From this standpoint, an officer should be more likely to shoot an armed, threatening person committing a robbery than a compliant, unarmed person pulled over for speeding. Nevertheless, in the existing literature the conditions that could alter the probability that an officer will use lethal force are either absent or relegated to control variables in analyses of racial differences. The tacit assumption is that the interaction stage is the same for both blacks and whites. Consequently, any racial differences in the decision to use lethal force are attributed to race itself, rather than potential racial differences in the circumstances leading up to that decision.

Examining the circumstances of police killings allows us to test this assumption and estimate the probability of racial differences in the interaction stage.⁴ Depending on the size

and direction of these differences, we can also estimate the likelihood of racial disparities in the contact and trigger stages as well. Beyond contributing to the literature on racial differences in lethal force, determining the source of the disparity is vital for identifying appropriate strategies to reduce the disproportionate risk to minorities as well as the total number of people killed by police.

For example, if I find that a lower proportion of black decedents threatened officers during the lethal interaction, the most likely explanation is that racial differences are occurring in the trigger stage. This would mean that the police are disproportionately killing unthreatening African Americans, consistent with a narrative of implicit racial bias. On the other hand, if I find that African American decedents threatened the police more often than white decedents, it is more probable that the disparity is occurring in the interaction stage.⁵ The implication of this scenario would be that the police are responding similarly to threats and that the racial disparity is likely due at least in part to the actions of black civilians.

However, if I find no systematic racial differences in the observable circumstances of police killings, the most likely explanation for the racial disparity in the rate of lethal force is that blacks are coming into contact with the police more frequently than whites. A null result would indicate that there are minimal racial differences in the interaction stage and at the same time negligible effects from bias in officers’ decisions to use lethal force. Although it is technically possible for there to be strong racial differences in these stages despite a null result, this would be implausible because it would require a perfect balance between the probability of events in the interaction stage and the frequency of racially biased decisions in the trigger stage.⁶ For this reason, I would interpret a null result as evidence that the racial disparity in lethal force is most likely driven by racial differences in police contact.

the inferences I can make about the possible source of the racial disparity in rates are more analytical and probabilistic rather than quantitative and causal.

5. An alternative explanation is that the level of demonstrated threat in the interaction stage is the same across race, but that police are biased against shooting threatening whites. However this is unlikely given research which finds that officers shoot white armed targets roughly the same amount as black armed targets (Correll et al. 2007).

6. For example, if I find no racial differences in the threat level of decedents but officers were more likely to shoot nonthreatening blacks than nonthreatening whites, the probability that an officer would shoot a nonthreatening black person must also equal the probability that a white person would threaten the police. In addition, the probability that the police would shoot a threatening white person would likewise need to equal the probability that a black person would not threaten the police.

3. However, even if there is not observable racial bias in officers’ decisions to use lethal force, racial bias could still be driving differences in the manner and rate at which officers come in contact with civilians (Knox, Lowe, and Mummolo 2019).

4. Because this analysis only examines cases where civilians died from lethal force, the scope and precision of these estimates is limited. Therefore

INTRODUCING THE POLICE KILLINGS AND PROTEST (PKAP) DATA SET

In order to determine whether and how the circumstances of police killings vary by race, I analyze an original data set of nearly 1,200 police-related deaths across the United States in 2015. The novel Police Killings and Protest (PKAP) data set builds upon the observations identified by the *Guardian* news group and greatly expands the available information to include hundreds of variables related to the lethal incident and collective action responses. In this paper, I analyze a subset of 126 of these variables, which detail the circumstances that led to the lethal interaction, what happened during encounter, as well as personal and demographic information about the decedent.

The PKAP data set is a major innovation because of both the breadth and depth of the police-related deaths it covers. Since the observations were identified through media reports, the data set contains over twice as many lethal incidents as federal sources that rely on voluntary reporting from police.⁷ For each police killing in the data set, my research team and I hand-coded detailed information from online news sources.⁸ Many of these news articles were referenced in the *Guardian* database, and the rest we identified by conducting internet searches using the name of the decedent and the location of their death as keywords. Each observation was double coded, and the author adjudicated any discrepancies between coders.

The majority of these online sources come from local newspaper or television reports. These news articles are generally based on official police statements and are highly consistent with police reports. Depending on local laws and department norms, officials may refrain from releasing certain information to the public; however, news reports regularly provide supplemental information based on interviews with witnesses or individuals close to the decedent. Consequently, these online sources regularly contain additional information unavailable in official police reports such as familial status, mental health history, or criminal record. The PKAP data set is therefore a taxonomy of virtually all the information the public could reasonably know about the decedent and the lethal interaction.

Though media reports are generally reliable and consistent sources of information, they may have their own racial biases that could affect the results of the analysis. For example, re-

search has shown that local news outlets tend to overemphasize crimes committed by African Americans (Entman and Rojecki 2001; Gilliam et al. 1996); consequently, journalists may decide to report criminal history differently for black and white decedents. Additionally, police departments themselves may release different types of information depending on the race of the decedent. For example, in the wake of the high-profile shooting in Ferguson, Missouri, police departments may feel a greater need to justify the shootings of African Americans and might release additional evidence in an attempt to stave off protest. Consequently, it is possible that the analysis could be biased toward finding racial differences in the circumstances of police killings even where none exist. Despite these potential limitations, the PKAP data set is ideal for investigating whether and how police killings vary by race.

RACE AS AN OUTCOME VARIABLE

In order to determine whether there are systematic differences in the circumstances of police killings across racial groups, I test whether it is possible to predict the race of decedents based on those circumstances. Although the police kill civilians of all ethnicities, due to empirical constraints I restrict the analysis to the two largest racial groups in the data: African Americans and non-Hispanic whites. The outcome variable for the analysis is a binary indicator of Race, where blacks are the reference category.

To predict race I analyze a random sample of 700 observations from the PKAP data set. In 2015, 51% of those killed by police were white and 27% were African American. Although blacks were killed at a rate three times that of whites, in absolute numbers twice as many whites were killed because of their share of the US population. Consequently, two-thirds of the decedents in the analysis sample are white (66%) and the other third are African American (34%).

FIVE CATEGORIES OF PREDICTORS

The analysis includes 126 predictors that capture the circumstances of the interaction and the characteristics of the decedent. In order to more succinctly discuss the variables and potential outcomes, I have grouped them into five distinct categories, summarized in table 1. These categories are based in part on existing variable groupings in the literature (Bolger 2015), but also on natural divisions in the data. The first set of predictors are Situational variables that detail the setting of the police-civilian interaction. Such variables include the location of the incident (residence, business, etc.), the time of day, how the police came in contact with the decedent (e.g., service call or traffic stop), and the presence of witnesses.

Variables in the Threat category assess the level of physical risk that the decedent posed to responding officers or nearby

7. The FBI counted 426 “justifiable homicides” by police in 2012, and the Bureau of Justice Statistics counted 497 police-related deaths in 2009. By contrast, the *Guardian* counted 1,146 police related deaths in 2015, and a report by the BJS suggests that this is a nearly complete accounting of all police killings in 2015 (Banks et al. 2016).

8. I would like to thank Evelin Hanhan, Atlanta Rydzik, Ben Streeter, and John Stempien for their excellent research assistance.

Table 1. Summary of Variable Categories

Category	Variable Examples	Number of Variables
Situational	Location: business Witness present Reason: service call Time of day: night	33
Threat	Unarmed Noncompliance: shoot first at officers Armed: knife Endangering civilians	26
Policing	Officer race: white SWAT Use chemical spray Use taser	32
Criminal activity	Crime: violent No crime in progress Criminal record Violent criminal record	12
Individual	Age Gender Married Mental health history	23
Total		126

civilians. This category includes information on whether the decedent possessed a weapon, the type of weapon, and whether they engaged in any noncompliant or aggressive behavior toward officers. Policing variables contain predictors related to the officers and their actions. These include indicators of whether a SWAT unit was deployed, whether the officer was undercover or off duty, and any nonlethal tactics the officers used before resorting to lethal force.

Predictors in the Criminal Activity category relate to a decedent's past or ongoing criminal involvement. This category includes indicators of any crimes that the decedent allegedly committed at the time of the police encounter as well as information about their criminal history. Finally, Individual variables include demographic and personal information about the decedent. These variables include the age and gender of the decedent as well as their mental health history and familial status (married, had children, etc.). With the exception of age, all of the variables included in the analysis are binary indicators that designate whether the particular attribute or characteristic was present in media reports. For brief descriptions and summary statistics of the variables, see ap-

pendix A1 (apps. A1–A3 are available online). All together, these five categories reveal a detailed picture of the observable context and circumstances of each lethal incident.

There is a possibility, however, that important features of the interaction could be absent from the data set. For example, officers have been shown to use less respectful language when speaking with African American drivers during traffic stops (Voigt et al. 2017). Features like language, tone, and nonverbal communication could be highly racially dependent in a lethal interaction but would not be captured in my data set. A further limitation is that for many variables I am not able to distinguish between missing data and a null value. For example, if a news article does not report that a decedent has children, it could be because they are not a parent or it could be because the media failed to report their familial status. Because of these possible data limitations, appropriate caution should be used in the interpretation of these results.

However, including such a wide range of predictors in the analysis increases our capacity to find racial differences in the data. For many of the included predictors, there are clear

existing theories regarding their potential relationship to race. For example, based on prior research African American decedents may be more likely to be unarmed or have a criminal history (DeGue, Fowler, and Calkins 2016; Shannon et al. 2017). Furthermore, any unobservable racial differences in the lethal interaction could directly or indirectly result in racial disparities in other aspects of the lethal interaction. Consequently, I include numerous additional variables related to the police killing so that the models are better able to detect any indirect or previously untheorized racial differences in the characteristics of police killings.

MACHINE LEARNING APPROACHES

Since the number of possible predictors is so large and many of the relationships between the variables, race, and lethal force have yet to be explored, this analysis is uniquely suited for machine learning techniques. Unlike traditional hypothesis testing methods, these techniques can detect complex relationships in the data and handle a large number of (possibly irrelevant) predictors. In this section I present the four machine learning techniques I use to attempt to predict the race of decedents and to identify which variables most distinguish between whites and blacks.

Random forest

A random forest is a supervised learning algorithm for classification, built from an ensemble of decision trees. Each decision tree can be represented as a series of if-then evaluations about the predictors (the branches), which then lead to conclusions about the classification outcome (the leaves or, in this case, race). Predictors which the tree places earlier in the evaluation process, closer to the “root” of the tree, are more important for distinguishing between classes. To prevent overfitting, each tree in the forest is created from a random sample of observations and predictors. When predicting the class of an observation, the random forest selects the modal or most popular prediction among all the decision trees. The accuracy measure I use for the random forest is derived from the “out-of-bag” error rate, the average proportion of incorrect predictions that the decision trees make on out-of-tree-sample observations. Because inferences about the likely source of racial disparities in lethal force depend on which variables are driving the results, I also present a separate random forest model for each of the five variable categories in addition to the full model.

Lasso regression

The second method I employ is lasso regression. Lasso is a model shrinkage and variable selection method for linear regression. Like ordinary least squares, lasso regression seeks

to minimize the sum of squared errors but with the added criteria of capping the absolute value of the sum of the coefficients. Within these constraints, the lasso model iteratively evaluates each predictor’s performance, setting to 0 coefficients that do not sufficiently minimize the error. Consequently, unlike a random forest, the lasso model includes only the subset of predictors that are most correlated with the outcome variable. With this measure I determine the accuracy rate using leave-one-out cross-validation and select the most parsimonious model that also minimizes the error.

Support vector machines and neural networks

In order to test the accuracy of two additional methods, support vector machines (SVMs) and neural networks. I divide the data set into a training and test sample. Rather than minimizing the error, the SVM seeks to maximize the distance between the binary categories across multidimensional space. The SVM uses a subset of the training points to construct a hyperplane that will separate blacks and whites across the many dimensions. Inspired by the structure of the human brain, neural network algorithms rely on a complex interaction of activation functions, hidden layers, and weights in much the same way that the brain relies on neurons, dendrites, and synapses for categorizing inputs.

RESULTS

To be predictive of race, the accuracy rate of the models must perform at least better than a naïve “no information” model that always predicts race to be “white.” Such a model would be accurate 66% of the time since that is the actual percentage of white decedents in the sample. Figure 1 presents the accuracy rates for the four machine learning techniques as well as the random forest models for each of the five variable categories.⁹

The main finding of this paper is that none of the models markedly outperform the no-information model, represented as a vertical line in the figure. The average accuracy rate for the full random forest model, lasso regression, SVM, and neural networks is only 67%, only marginally better than that of the no-information model. The best-performing method is the random forest, which has a cross-validation accuracy rate of 69%. However, this gain is fairly marginal and indicates that the many variables included in the data set have only minimal collective racial differences. These null results are robust across machine learning methods as well as alternate classification cutoffs.¹⁰

9. See app. A2 for more detailed model outputs.

10. The results are equivalent if I change the prediction cutoff value from the standard .5 to .66 to accommodate the unbalanced categories.




Figure 1. Machine learning accuracy rates. This graph presents the prediction accuracy rate for different machine learning methods. Higher values indicate that the method is better able to distinguish between racial groups. The vertical line represents the accuracy rate of a no-information model, which always predicts race to be white. “RF” is an abbreviation for random forest.

The results for the category-based models do reveal that some variable groups are more predictive of race than others. The random forest model for the Individual category contains only 23 variables but performs better than the no-information model. Because the other four category-specific models perform worse than the no-information model, the marginal racial differences in the data are likely concentrated within the Individual category. In the next section, I discuss some of these specific variables in more detail and explain what the results mean for our understanding of racial differences across the three stages of the lethal interaction process.

DISCUSSION

This low prediction accuracy indicates that there are likely only marginal racial differences in the observable circumstances of police killings. Although African Americans are more likely to be killed by police, as a whole their lethal interactions are remarkably similar to those of white decedents.¹¹ These models are capable of detecting complex dependencies among the predictors and nonlinear relationships with the outcome variable—yet they find no aggregate patterns that are strongly predictive of race. Though a few individual variables may have statistically significant differences in means, the null results suggest that the combined racial differences across all of the variables are not sufficiently large to even

partially explain the racial disparity in the rate of police killings.

On the other hand, these results do indirectly point to the probable source of the threefold racial difference in the rate of police killings. As previously discussed, the null result indicates that the racial disparity is most likely originating from the contact stage of the lethal interaction process. Because the analysis only examines incidents where a person was killed, it is not possible to confirm this conclusion quantitatively, but analytically it is improbable for there to be strong racial differences in either the interaction or trigger stage given the findings. As a result the most likely explanation for the difference in rates is that African Americans have more frequent contact with the police, leading to a higher risk of being killed by law enforcement.

What is driving this higher rate of minority contact with law enforcement? Empirical studies on the mechanics of police contact focus almost exclusively on racial differences in police-initiated interactions such as traffic and pedestrian stops (Neil and Winship 2019). However, according to the PKAP data set only 11% of police killings were initiated from these kinds of stops. Instead the majority of deaths resulted from civilian-initiated service calls. Though implicit bias may certainly be driving disparate rates in police-initiated contact, the mechanisms by which racial bias could be causing higher contact rates are less obvious for non-officer-initiated contact. To develop some intuitions for these mechanisms, we can examine the most important variables in the machine learning models.

11. Nix et al. (2017) also find no difference between blacks and white decedents in terms of the frequency with which they attacked officers or civilians.

Both lasso regression and random forests provide clear criteria for evaluating which features are most important for predicting race. For the lasso regression, the important variables are those included in the pruned model, which in this case includes 28 of the 126 variables. The random forest retains all the variables but ranks variable importance based on how much a predictor increases the average accuracy of decision trees. To determine the most influential variables overall, I subset the top 10 most important predictors for each model and identified which ones were overlapping. The models have a good deal of agreement, and the five overlapping variables stand out as most predictive of race.¹²

Of these five top variables, three belong to the Individual category, explaining the high performance of that category-specific model. Interestingly, all three variables are more correlated with white decedents. Whites killed by police are more likely to be suicidal at the time of their death and to have a history of mental health issues. At the same time, they are also more frequently reported as being intoxicated during the lethal interaction and as having a history of substance abuse problems. Consequently, these findings suggest that compared to blacks, whites may be coming in contact with the police more frequently as the result of acute, and sometimes chronic, emotional crises.

By contrast, variables that most distinguish the deaths of African Americans are related to officer characteristics. Black decedents are more likely to have been killed by an officer who was either undercover or off duty as well as by black officers. Though officers' race is rarely provided in news reports, this finding underscores potential differences in how black officers are deployed across neighborhoods. The higher rate of lethal contact with plainclothes officers suggests that some black decedents may not have been aware that they were engaging with law enforcement until late in the lethal interaction. However, because only a small proportion of black decedents were killed under these circumstances, the predominant source for the racial disparity in police contact remains unidentified.

Even more informative are the predictors missing from the top variable lists. Based on the racial similarities and lack of predictive power for variables in the situational and criminal categories, the higher rate of lethal police contact for blacks is likely not driven by existing explanations related to discretionary stops or decedents' alleged criminal activity (Fridell 2016; Goff et al. 2016). Though the top Individual variables suggest that there may be a unique profile for some white decedents, the results do not provide clear reasons for

the higher rate of black contact with law enforcement. In order to fully understand why blacks are disproportionately killed, further research is needed to identify why African Americans are coming in greater contact with the police through so many different means and under such a wide array of circumstances.

CONCLUSION

This paper demonstrates that, as a whole, the police are killing white and black decedents under very similar circumstances. Even with advanced machine learning techniques and the high-dimensional PKAP data set, it is not possible to predict decedents' race more accurately than naïve assignment. These null results suggest that the threefold racial difference in the rate of police killings is most likely due to disproportionate African American contact with the police, rather than racial differences in the circumstances of the interaction or bias in officers' decision to use lethal force. Though these findings lend support to the contact-based explanations in the literature, including those which point to racial bias in the initiation of police contact, the variable importance results also indicate that further research is needed to identify why the disproportionate contact is occurring under such a wide array of circumstances.

I motivated this paper with a father's assertion that if his son had been white, the police would not have shot him. The findings of this paper demonstrate the partial truth of this statement. On one hand, the results indicate that if a white person encountered officers under similar circumstances, they would likely be killed by police. However, if we take into consideration the probability of such an encounter's occurring in the first place, the risk that a white person would be killed by police is small compared to that of a black person.

The reason why so many police killings of African Americans have sparked outrage is that, at least to many, the circumstances of those interactions did not appear to warrant lethal force. A jarring implication of my research is that an analogous proportion of white decedents are also killed by police under similarly dubious circumstances. This finding in no way decreases the egregiousness of the disproportionate killings of African Americans but, rather, increases the urgency of reducing police killings across all racial groups.

ACKNOWLEDGMENTS

I would like to thank Kirk Bansak, Lindsay Hundley, Hakeem Jefferson, David Laitin, and Dan Thompson for their insightful comments. I also thank the attendees of the Race and Law Enforcement Symposium at the University of Southern California as well as the anonymous reviewer for their feedback. A special thank you to Edgar Franco Vivanco for his support and advice throughout the project.

12. See app. A3 for full variable importance results.

REFERENCES

- Banks, Duren, Paul Ruddle, Erin Kennedy, and Michael G. Planty. 2016. "Arrest-Related Deaths Program Redesign Study, 2015–16: Preliminary Findings." Technical report, US Department of Justice, Bureau of Justice Statistics, Washington, DC.
- Baumgartner, Frank R., Leah Christiani, Derek A. Epp, Kevin Roach, and Kelsey Shoub. 2017. "Racial Disparities in Traffic Stop Outcomes." *Duke Forum for Law and Social Change* 9 (21): 21–53.
- Bolger, P. Colin. 2015. "Just Following Orders: A Meta-analysis of the Correlates of American Police Officer Use of Force Decisions." *American Journal of Criminal Justice* 40 (3): 466–92. <https://doi.org/10.1007/s12103-014-9278-y>.
- Buehler, James W. 2017. "Racial/Ethnic Disparities in the Use of Lethal Force by US Police, 2010–2014." *American Journal of Public Health* 107 (2): 295–97. <https://doi.org/10.2105/AJPH.2016.303575>.
- Chan, Melissa. 2016. "Minnesota Governor: Philando Castile Would Be Alive If White." *Time*, July 7. <http://time.com/4397248/minnesota-gov-mark-dayton-philando-castile-shooting/>.
- Correll, Joshua, Bernadette Park, Charles M. Judd, and Bernd Wittenbrink. 2002. "The Police Officer's Dilemma: Using Ethnicity to Disambiguate Potentially Threatening Individuals." *Journal of Personality and Social Psychology* 83:1314–29. <https://doi.org/10.1037/0022-3514.83.6.1314>.
- Correll, Joshua, Bernadette Park, Charles M. Judd, Bernd Wittenbrink, Melody S. Sadler, and Tracie Keese. 2007. "Across the Thin Blue Line: Police Officers and Racial Bias in the Decision to Shoot." *Journal of Personality and Social Psychology* 92 (6): 1006–23. <https://doi.org/10.1037/0022-3514.92.6.1006>.
- DeGue, Sarah, Katherine A. Fowler, and Cynthia Calkins. 2016. "Deaths Due to Use of Lethal Force by Law Enforcement: Findings from the National Violent Death Reporting System, 17 U.S. States, 2009–2012." *American Journal of Preventive Medicine* 51 (5S1): 173–87. <https://doi.org/10.1016/j.amepre.2016.08.027>.
- Eberhardt, Jennifer L., Phillip Atiba Goff, Valerie J. Purdie, and Paul G. Davies. 2004. "Seeing Black: Race, Crime, and Visual Processing." *Journal of Personality and Social Psychology* 87:876–93. <https://doi.org/10.1037/0022-3514.87.6.876>.
- Eckhouse, Laurel. 2018. "Everyday Risk Disparate Exposure and Racial Inequalities in Police Violence." Unpublished manuscript.
- Entman, Robert M., and Andrew Rojecki. 2001. *The Black Image in the White Mind: Media and Race in America*. Chicago: University of Chicago Press.
- Fridell, Lorie A. 2016. "Explaining the Disparity in Results across Studies Assessing Racial Disparity in Police Use of Force: A Research Note." *American Journal of Criminal Justice* 42 (3): 502–13. <https://doi.org/10.1007/s12103-016-9378-y>.
- Fryer, Roland G. 2019. "An Empirical Analysis of Racial Differences in Police Use of Force." *Journal of Political Economy* 127 (3): 1210–61.
- Gelman, Andrew, Jeffrey Fagan, and Alex Kiss. 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): 813–23. <https://doi.org/10.1198/016214506000001040>.
- Gilliam, Franklin D., Shanto Iyengar, Adam Simon, and Oliver Wright. 1996. "Crime in Black and White: The Violent, Scary World of Local News." *Harvard International Journal of Press/Politics* 1 (3): 6–23.
- Goff, Phillip Atiba, John Jay, 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, Los Angeles.
- Guardian. 2017. "Police Officer Who Shot Dead Philando Castile Acquitted of All Charges." From Associated Press, St. Paul, MN, June 16. <https://www.theguardian.com/us-news/2017/jun/16/philando-castile-death-police-officer-not-guilty>.
- Holmes, Malcolm D., Brad W. Smith, Adrienne B. Freng, and Ed A. Muñoz. 2008. "Minority Threat, Crime Control, and Police Resource Allocation in the Southwestern United States." *Crime and Delinquency* 54 (1): 128–52. <https://doi.org/10.1177/0011128707309718>.
- Knox, Dean, Will Lowe, and Jonathan Mummolo. 2019. "The Bias Is Built In: How Administrative Records Mask Racially Biased Policing." *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3336338>.
- Lerman, Amy E., and Vesla M. Weaver. 2014. "Staying out of Sight? Concentrated Policing and Local Political Action." *ANNALS of the American Academy of Political and Social Science* 651:202–19.
- Lundman, Richard J., and Robert L. Kaufman. 2003. "Driving While Black: Effects of Race, Ethnicity, and Gender on Citizen Self-Reports of Traffic Stops and Police Actions." *Criminology* 41 (1): 195–220.
- Meehan, Albert J., and Michael C. Ponder. 2002. "Race and Place: The Ecology of Racial Profiling African American Motorists." *Justice Quarterly* 19 (3): 399.
- Miller, Ted R., Bruce A. Lawrence, Nancy N. Carlson, Delia Hendrie, Sean Randall, Ian R. H. Rockett, and Rebecca S. Spicer. 2017. "Perils of Police Action: A Cautionary Tale from US Data Sets." *Injury Prevention* 23 (1): 27–32. <https://doi.org/10.1136/injuryprev-2016-042023>.
- Mummolo, Jonathan. 2018. "Militarization Fails to Enhance Police Safety or Reduce Crime but May Harm Police Reputation." *Proceedings of the National Academy of Sciences* 115 (37): 9181–86. <https://doi.org/10.1073/PNAS.1805161115>.
- Neil, Roland, and Christopher Winship. 2019. "Methodological Challenges and Opportunities in Testing for Racial Discrimination in Policing." *Annual Review of Criminology* 2 (1): 73–98. <https://doi.org/10.1146/annurev-criminol-011518-024731>.
- Nix, Justin, Bradley A. Campbell, and Edward H. Byers. 2017. "A Bird's Eye View of Civilians Killed by Police in 2015: Further Evidence of Implicit Bias." *Criminology and Public Policy* 16 (1). <https://doi.org/10.1111/1745-9133.12269>.
- Shah, Khushbu. 2018. "If He Had Been White They Wouldn't Have Shot Him." *Guardian*, December 9. www.theguardian.com/us-news/2018/dec/09/ej-ematic-fitzgerald-bradford-jr-shooting-alabama-mall-family.
- Shannon, Sarah K. S., Christopher Uggen, Jason Schnittker, Melissa Thompson, Sara Wakefield, and Michael Massoglia. 2017. "The Growth, Scope, and Spatial Distribution of People with Felony Records in the United States, 1948–2010." *Demography* 54:1795–1818. <https://doi.org/10.1007/s13524-017-0611-1>.
- Stults, Brian, and Eric P. Baumer. 2007. "Racial Context and Police Force Size: Evaluating the Empirical Validity of the Minority Threat Perspective." *American Journal of Sociology* 113 (2): 507–46. <https://doi.org/10.1086/518906>.
- Swaine, Jon. 2014. "Father of Ohio Man Killed by Police: My Son Might Be Alive If He Were White." *Guardian*, September 25. <https://www.theguardian.com/world/2014/sep/25/father-ohio-man-walmart-killed-police-john-crawford>.
- Todd, Andrew R., Kelsey C. Thiem, and Rebecca Neel. 2016. "Does Seeing Faces of Young Black Boys Facilitate the Identification of Threatening Stimuli?" *Psychological Science* 27 (3): 384–93. <https://doi.org/10.1177/0956797615624492>.
- Voigt, Rob, Nicholas P. Camp, Vinodkumar Prabhakaran, William L. Hamilton, Rebecca C. Hetey, Camilla M. Griffiths, David Jurgens, Dan Jurafsky, and Jennifer L. Eberhardt. 2017. "Language from Police Body Camera Footage Shows Racial Disparities in Officer Respect." *Proceedings of the National Academy of Sciences* 14 (25): 6521–26. <https://doi.org/10.1073/pnas.1702413114>.
- Worrall, John L., Stephen A. Bishopp, Scott C. Zinser, Andrew P. Wheeler, and Scott W. Phillips. 2018. "Exploring Bias in Police Shooting Decisions with Real Shoot/Don't Shoot Cases." *Crime and Delinquency* 64:1171–92. <https://doi.org/10.1177/0011128718756038>.