Reply to Schimmack: Critique of Johnson et al. (2019)

David J. Johnson

University of Maryland at College Park
Michigan State University

Joseph Cesario

Michigan State University

Recently we published a paper showing that officer race did not relate to the race of a civilian shot and, additionally, a lack of anti-Black racial disparities among those fatally shot by police (Johnson, Tress, Burkel, Taylor, & Cesario, 2019). In response, Schimmack (2019) raised a number of criticisms, both on his blog and directly to us by email. In his view, these criticisms undermine the research so fundamentally that the paper should be retracted.

In this reply, we address each of the points he raised. These include: (1) coding of variables; (2) interpretations of model intercepts; (3) whether we considered population proportions and crime rates in our analyses; (4) the nature of the age effect and whether this reveals racial bias; and (5) additional analyses by Schimmack claiming to show evidence of racial disparities.

As we show below, these criticisms are either inaccurate or do not change the substantive conclusions of the paper. In the one instance where the critique (Critique I) does change the size of our effects (but not their direction or significance), we are in contact with *PNAS* to determine how to add these additional analyses to the record of the paper.

Critique I: Coding of Variables

The first criticism concerns the coding scheme used in our analyses. To be clear up front, the implications of this criticism impact the *size* of our estimates of racial disparities, not their direction or our substantive conclusions. The critique centers on our tests of whether a person fatally shot is more likely to be Black (or Hispanic) than White when crime rates are equal across race. These tests are reported in Table 3 of the original manuscript. In these intercept tests we centered all other continuous variables and coded categorical variables with unweighted effect coding. For instance, unarmed shootings were coded as $-.5$ and armed shootings as $.5$.

As Schimmack correctly points out, the intercept values reported in Table 3 of the original manuscript need to be interpreted with caution because the number of cases across

coding categories are unequal (te Grotenhuis et al., 2017). For example, fatal shootings were more common among citizens who were not suicidal or who did not have mental health issues. When using unweighted effect coding the value of the intercept will be skewed in the direction opposite the more frequently coded occurrence (i.e., when civilians are suicidal and have a mental health issue).

As White individuals were more likely to be suicidal and have mental health issues than Black and Hispanic civilians, unweighted effect coding may underestimate the likelihood that a person fatally shot is Black or Hispanic. Weighted effect coding provides a more liberal test of racial disparity given higher rates of mental health issues and suicide risk among White civilians fatally shot. This is important as these model intercepts are the key test of anti-Black or anti-Hispanic disparity amongst those fatally shot by police when controlling for racial differences in crime rates. Accordingly, we reran the analyses reported in Table 3 of (Johnson et al., 2019) using weighted effect coding. These results are reported in Table 1 below.

Consistent with the results in our published paper, when using weighted effect coding and controlling for differences in crime rates, *overall and in no specific type of shooting did we find significant evidence of anti-Black or anti-Hispanic disparity among those fatally shot*. Thus, this substantive point has not changed from the original paper. However, the overall odds ratios are less extreme than the strong anti-White disparity originally reported. When controlling for racial differences in crime rates, in the average shooting (Model 0), a person fatally shot was 3.2 times more likely to be White than Black, and 2.8 times more likely to be White than Hispanic (compared to 6.7 and 3.3, respectively, when using unweighted effect coding).

In sum, even when using weighted effect coding, the conclusion that we do not observe evidence of anti-Black or anti-Hispanic disparity across different types of shootings still holds. That said, we agree that weighted effect coding provides a more balanced estimate of the magnitude of these effects than unweighted effect coding. We thank Schimmack for raising this issue to our attention.

Critique II. Interpretation of the Intercept

A second claim by Schimmack is that our analyses must be fundamentally flawed because the intercepts of our regres-

David Johnson is postdoctoral researcher in the Lab for Applied Social Science Research at the University of Maryland. Joseph Cesario is associate professor of psychology at Michigan State University. The original Johnson et al. (2019) manuscript and the Schimmack (2019) critique are available at the second author's website.

Table 1
Racial Disparity in Civilian Race by Shooting Type

Model and Variable	Level	Black		Hispanic	
		OR	95% CI	OR	95% CI
0. Typical Shooting	–	0.31	0.23, 0.42	0.36	0.28, 0.47
1. Civilian Age	Low	0.54	0.37, 0.78	0.66	0.47, 0.92
2. Civilian Age	High	0.18	0.13, 0.26	0.20	0.14, 0.29
3. Civilian Armed	No	0.52	0.25, 1.09	0.32	0.13, 0.77
4. Civilian Armed	Yes	0.29	0.21, 0.41	0.37	0.27, 0.49
5. Civilian MH issue	No	0.37	0.26, 0.51	0.45	0.33, 0.63
6. Civilian MH issue	Yes	0.19	0.11, 0.32	0.18	0.10, 0.35
7. Civilian suicidal	No	0.36	0.26, 0.49	0.35	0.26, 0.48
8. Civilian suicidal	Yes	0.10	0.05, 0.22	0.44	0.17, 1.14
9. Civilian attacking	No	0.16	0.04, 0.55	0.37	0.12, 1.11
10. Civilian attacking	Yes	0.33	0.24, 0.44	0.36	0.28, 0.48
11. Officer number	Low	0.30	0.21, 0.43	0.33	0.23, 0.46
12. Officer number	High	0.32	0.22, 0.47	0.40	0.29, 0.56
13. Officer % Black	Low	0.29	0.20, 0.43	0.35	0.24, 0.52
14. Officer % Black	High	0.33	0.24, 0.47	0.37	0.27, 0.52
15. Officer % Hispanic	Low	0.25	0.17, 0.37	0.29	0.20, 0.41
16. Officer % Hispanic	High	0.38	0.27, 0.56	0.46	0.34, 0.62
17. Officer % women	Low	0.30	0.21, 0.43	0.38	0.27, 0.56
18. Officer % women	High	0.33	0.23, 0.46	0.34	0.24, 0.48
19. Officer experience	Low	0.30	0.21, 0.43	0.36	0.25, 0.51
20. Officer experience	High	0.33	0.23, 0.46	0.37	0.25, 0.53

Note. All variables are weighted effects coded or centered with the exception of the difference score between White and Black (or White and Hispanic) crime rates. This was done to control for racial differences in violent crime rates (see the Supplemental Materials of Johnson et al., 2019). Models 1 - 20 are coded to indicate certain types of shootings. Level indicates what level of the variable racial disparity is tested. MH = mental health.

sion models (when converted to odds ratios) deviate from the marginal odds ratio estimated from the observed data. For example, he states:

“It makes no sense when the authors claim that they centered continuous variables and effect coded nominal variables and the intercept changed from $\exp(-.71) = .49$ to $\exp(-1.90) = .15$, which they report in Table 3. Something went wrong in their analyses.”

The claim that we made an error in our analysis is unfounded because Schimmack’s claim that the intercept should not change is not correct. We first discuss conceptually why Schimmack’s claim is incorrect, and then falsify it in two ways: 1) by analyzing the data from our PNAS paper and 2) by using a simple simulation.

Conceptual Problems

In logistic (and multinomial) regression, the intercept refers to the log odds of the outcome when all predictors are

at zero. In our paper, the outcome is civilian race, which has three categories: Black, Hispanic, or White; White is used as a reference group. Thus, there are two intercepts in our models, each of which indicates the log odds that a person fatally shot is Black (or Hispanic) relative to White.

Schimmack’s assertion is that our model intercepts—when expressed as an odds ratio—should not change from the marginal odds that a person fatally shot is Black or Hispanic *when all variables are centered*. He asserts that the intercept should “represent the relative frequencies of Black/White in the sample (with some small deviation that I still don’t fully understand but that is irrelevant).”

This claim is demonstrably false. The intercept from a logistic regression is only guaranteed to mirror the relative frequencies in the sample when the model contains no predictors. When predictors are included, the intercept will change. This is because the intercept refers to the log odds of the outcome when all predictors are at zero (Menard, 2002). When predictors are centered, the intercept represents the log odds of the outcome when all predictors are at their mean.

This value can and does deviate from the relative frequencies (marginal odds) of the sample.

Falsification with Sample Data

The central claim in this argument is that the value of the intercept should not change when all predictors are centered from the relative frequencies in the sample. In our sample, 917 people were fatally shot: 501 Whites, 245 Blacks, and 171 Hispanics. Thus the relative frequency (i.e., marginal odds) of a Black person being shot relative to a White person is $(245/501) = 0.49$. The marginal odds of a Hispanic person being shot relative to a White person is $(171/501) = 0.34$.

Schimmack's claim is that these intercept values should not change from .49 and .34, so long as all predictors are centered (whether continuous or categorical). We show that this is false by running two simplified versions of our analyses using the publicly available data from our paper. The data and R script for this analysis can be found at <https://osf.io/cp6hf/>.

Specifically, mirroring Schimmack's critique, we ran two logistic regressions, 1) predicting the odds of being Black (relative to White), and 2) the odds of being Hispanic (relative to White). In the analysis of Black civilians, Hispanics were removed from the sample, and vice versa.

First, we verified that the intercept is identical to the relative frequencies in the sample when the model contains no predictors, i.e., is a null model. These results are reported in Table 2 below, and show that the model intercept is identical to the marginal odds in the special case where the regression model contains no additional predictors.

Table 2
Intercept Changes Across Models

Model	OR (Black)	OR (Hispanic)
Marginal Odds	0.49	0.34
Intercept Only	0.49	0.34
% Homicide	0.39	0.25

Note. Marginal odds were calculated from relative frequencies. % Homicide is race-specific.

Next, we ran a model that included a single centered predictor, race-specific homicide percentages (% Black homicide for the logistic regression of Black vs. White, and % Hispanic homicide for the logistic regression of Black vs. White). As shown in Table 2 the model intercept changes. The intercept for Black civilians changes from 0.49 to 0.39, and the intercept for Hispanic civilians changes from 0.34 to 0.25. Including more predictors (e.g., White crime rates, income inequality) further decreases in the intercept (when converted to an odds ratio), even when those predictors are centered. Thus, Schimmack's claim that the intercept should not change is not supported by a reanalysis of our publicly available data.

The reason that our model-predicted odds ratios differ from the marginal odds is because we are estimating the odds of being Black (or Hispanic) vs. White when predictors are at their mean. This is exactly why the test of the intercept is useful for testing racial disparities—it allows us to test racial disparities in the typical (i.e., average) shooting.

Falsification with Simulated Data

The claim that a logistic regression intercept should always reflect the marginal odds when all predictors are centered can be further falsified by a simple simulation analysis. The reproducible R script for this simulation can be found at <https://osf.io/cp6hf/>.

We simulated a dataset of 100 trials. The probability of success (i.e., a 1 vs. a 0) on any given trial was estimated using the following formula:

$$\log\left(\frac{p}{1-p}\right) = \beta_0 + \beta_1 X \quad (1)$$

Where β_0 is the intercept and was set at -0.5, and β_1 controls the strength of the association between the predictor X and success and was set to 0.05. The predictor X was set to uniformly vary from 0 to 100 ($M = 53$).

In our simulated dataset, the marginal odds of success were $(77/23) = 3.3$. Thus, a value of 1 was observed 3.3 times more than a value of 0. A logistic regression with the intercept only also produced an intercept that (converted to an odds ratio) is 3.3. Again, in the special case where the regression model contains no additional predictors, the model intercept is identical to the marginal odds (3.3). However, when we add the centered predictor X into the analysis, the intercept changes from 3.3 to 6.4, indicating that the odds of a 1 (vs. a 0) are 6.4 when X is at its average. Again, Schimmack's claim that the intercept should not change is not supported by a simulation approach.

The script enables users to change the values in the simulation to verify that these conclusions are robust to how the model is specified. It is worth noting that the model intercept increasingly diverges from the marginal odds as the association between the predictor X and success increases. Said differently, the model intercept will reflect the marginal odds when a predictor is present *only* if the predictor is unrelated to the outcome variable, i.e., $\beta_1 = 0$. This is clearly not the case in our analysis, where race-specific crime rates strongly predict the race of a person fatally shot.

Critique 3: Taking Population and Crime into Account

Schimmack at various points claims that we failed to take population rates into account, which prevent clear interpretation of our results. For example, on his blog post, he states:

“The authors have an unscientific aversion to take population rates into account. In counties

with mostly White population, crime is mostly committed by White citizens, and police is more likely to encounter and kill White criminals. It is not a mistake to include population rates in statistical analyses. It is a mistake not to do so.”

and:

“The interpretation of this result as a measure of racial disparities is wrong. One factor that is omitted from the analysis is the proportion of White citizens in the counties. It doesn’t take a rocket scientist to realize that counties with a larger White population are more likely to have White victims. The authors do not take this simply fact into account, although they did have a measure of the population size in their data set.”

Schimmack is correct in noting that different counties will have different proportions of Black and White citizens. As reported the Supplemental Materials, county crime rates and population rates both vary by race in our sample (see Table 3). Without correcting for how often people of different rates are exposed to police, the fact that more civilians in our sample were White (501) versus Black (245) or Hispanic (171) does not test whether or not a Black or Hispanic person is more *likely* to be shot than a White person.

Table 3
Crime and Population Rates

Race	Homicide	Population
White	53%	68%
Black	28%	11%
Hispanic	14%	16%

However, as we discuss in the Supplemental Materials, population differences are a poor proxy for understanding disparities in the *likelihood* of being fatally shot. The ideal proxy would be racial differences in exposure to police in situations where fatal force is likely. While such exact records do not exist on a national level, in the paper we use homicide rates as a closer proxy for exposure to police, given that the vast majority of fatal shootings take place in the context of violent crime (Cesario, Johnson, & Terrill, 2019; Nix, Campbell, Byers, & Alpert, 2017).

Where Schimmack is unequivocally incorrect is his statement that we failed to account for different percentages of either Black and White citizens or Black and White crime rates within counties. In the main text, our tests of the intercept (Table 3; Johnson et al., 2019) control for differences in crime rates, which we argue are a better proxy for police exposure. However, we know that this is an area of debate. Therefore, in the Supplemental Materials we clearly test for racial disparities controlling for differences in crime rates

alone, population rates alone, as well as crime and population rates simultaneously.

As our use of unweighted effect coding affects the tests of the intercept reported in the Supplemental Materials. Therefore, we reanalyzed the racial disparities reported in the section “Additional Tests of Racial Disparities” using weighted effect coding. These results are reported in Table 4 below.

When we use weighted effect coding and control for crime rates alone, we do not find evidence of anti-Black disparity—anti-White disparity is actually significant. If we control for population rates alone, we obtain the opposite result: a person fatally shot is 2.0 times more likely to be Black than White. These opposing findings reflect the effects of population vs. crime rates as a measure for police exposure. It is widely known that anti-Black disparity is observed when population rates are used as a benchmark (Geller & Karales, 1981; Hehman, Flake, & Calanchini, 2018; Jacobs, 1998; Ross, 2015; Scott, Ma, Sadler, & Correll, 2017; Takagi, 1974), whereas when violent crime is used as a benchmark, anti-Black disparities in fatal officer-involved shootings (FOIS) disappear or reverse (Cesario et al., 2019; Fyfe, 1982; Goff, Lloyd, Geller, Raphael, & Glaser, 2016; Inn, Wheeler, & Sparling, 1977).

As we argue in the main paper, using population as a benchmark makes the strong assumption that White and Black civilians have equal exposure to situations that are likely to result in FOIS. This assumption is unlikely to be true (as shown in Table 3), which is why we focused on using violent crime as a benchmark, as the majority of FOIS involve armed civilians Nix et al. (2017). However, a benefit of our approach is that we can control for population and crime differences simultaneously. *When we do this we do not find evidence of anti-Black disparity.* A person fatally shot is 1.1 times more likely to be White than Black, and this disparity is not significant.

The point of these additional tests was to serve as a sensitivity analysis to examine the robustness of our conclusions based on different proxies for exposure. When using weighted effect coding we can see that the nature of these conclusions depends on what proxy for police exposure is controlled for. What is clear is that we did not have an “unscientific aversion” to use population rates because we in fact conducted such analyses in the Supplemental Materials.

Critique 4: Age and Race Bias

Another criticism by Schimmack is that we *did* find evidence of racial disparity but failed to note it or appreciate it fully. This claim is based on a significant effect of age in Table 2, demonstrating that as age of the civilian increased, odds were greater that the civilian was White vs. Black; said equally, as age of the civilian decreased, odds were greater that the civilian was Black vs. White. As we stated in the article, “Older civilians were 1.85 times less likely (OR =

Table 4
Racial Disparity in Civilian Race

Model	Intercept Controls For	Black		Hispanic	
		OR	95% CI	OR	95% CI
S0	Crime Differences	0.31	0.23, 0.42	0.36	0.28, 0.47
S1	Population Differences	2.03	1.21 3.41	0.30	0.14, 0.64
S2	Crime and Population Differences	0.89	0.44, 1.80	0.32	0.12, 0.81

Note. Differences were controlled for by creating a score indicating the difference between White and Black (or Hispanic) crime and population in a given county. When this uncentered predictor is entered into the model, the intercept tests for racial disparity in a county where crime rates, population rates, or both do not vary by race.

0.54 [0.45, 0.66]) to be Black than White.”

As Schimmack notes:

“Thus, the authors found that young Black citizens are killed more frequently by police than young White men, even when they attempted to control statistically for the fact that young Black men are disproportionately involved in criminal activities.”

and:

“Third, the authors ignored a key finding of their own analysis that age is a significant predictor of police shootings. Consistent with the Black Lives Matter claim, their data show that police disproportionately shoots young Black men.”

Some clarification appears to be in order about the age effect as shown in the original Table 2. We do not know for sure that “young Black citizens are killed more frequently by police than young White men” based on the significant age effect. This age effect could be driven by young Black men being fatally shot at disproportionate rates vs. young White men (as Schimmack suggests). It also could be that older White men are fatally shot more than older Black men, with little race difference among young men. For instance, at young ages there may be no race effect while at older ages, a greater proportion of older White men (e.g., with mental health problems) may increase the anti-White disparity at older ages. It could even be that there is anti-White disparity for *both* young and old civilians but that this anti-White disparity is *smaller* in magnitude for young civilians. All these patterns would produce a significant age coefficient. We do not say this to dismiss the real possibility of anti-Black disparity, but to stress the limitations about what conclusions can be drawn from this model coefficient.

For some evidence relevant to the nature of this effect, consider the analyses presented in Table 1 (above). Looking at typical shootings for young civilians versus old civilians (Models 1 and 2), we do find that the odds are relatively more anti-Black in the shootings of young civilians

(0.54) compared to shootings of older civilians (0.18). However, *in both cases there is no significant anti-Black disparity*, i.e., the odds ratios are significantly less than 1.0. These results merely show a pattern of *less* anti-White disparity in the cases of young civilians. This does not, as Schimmack suggests, present clear evidence of anti-Black disparity, which is why we chose to not emphasize this age effect in the paper and why we described our results as “not providing evidence of anti-Black disparity.” When controlling for differences in crime rate, there is not significant anti-Black disparity for younger or older civilians.

Critique 5: Additional Analyses

On his blog post, Schimmack presents some partial re-analyses of our data which reveal anti-Black disparities. Important differences between the analytic decisions in our paper and the analytic decisions in Schimmack’s critique impact the crucial test of the intercept value. Our approach to controlling for exposure was a difference approach adopted from the analyses in Scott et al. (2017) and differs slightly from the ratio approach used in the critique. The approach in the critique also removes over 100 cases from the dataset, which makes the datasets less comparable. The critique does not control for the fact that shootings are nested within counties. And most importantly, the critique only includes a small subset of the predictors in the model, which greatly impacts the test of the model intercept.

We do not just point out these differences to simply criticize Schimmack’s approach. We note these differences because they impact the intercept test for racial disparity amongst those fatally shot. We justify our decisions in the paper (e.g., controlling for crime rates vs. population rates) based on past work, but know that researchers may disagree with these decisions; the issue of benchmarking and exposure has been debated in criminal justice for decades.

Reasonable researchers can debate our methodological and analytic choices because we made these public and present alternative choices for the reader to consider. As clarified above, the differences between our analyses and Schim-

mack's are due to different methodological choices and his use of a subset of the data, not because of errors.

Concluding Comments

We thank Schimmack (2019) for his critique of our paper. Post-publication peer review such as this plays an important role in improving published findings and advancing research generally. In particular, Schimmack's suggestion to use weighted effects coded leads to more balanced tests of racial disparity given that the frequencies of our categorical variables are unbalanced. However, we stress that although this suggestion impacts the magnitude of our estimates of racial disparity, it does *not* change our overall substantive conclusion—we did not observe anti-Black or anti-Hispanic disparity amongst those fatally shot when controlling for racial differences in crime rates.

Although post-publication peer review can improve published findings and point out flaws that may have been missed during the review process, this does not always mean that the critique will be without flaws. For example, we demonstrate that Schimmack's misunderstanding of the intercept in a logistic regression led him to the erroneous conclusion that changes in our intercept from the marginal odds must have been due to an error. Similarly, his claim that we did not control for racial differences in population proportions is clearly false as we did exactly this in the Supplemental Materials.

The issue of racial disparities in fatal shootings is a hotly debated issue that is difficult to study because of limited data on policing. As a result, we employed a novel method to test these disparities and did not observe anti-Black or anti-Hispanic disparity amongst those fatally shot. Others may disagree with some of our methodological choices, such as controlling for racial disparities using crime or population rates. Because our data is publicly available and post-publication peer review is transparent, others can see the debate over these choices and ultimately decide for themselves how best to analyze the data.

References

- Cesario, J., Johnson, D. J., & Terrill, W. (2019). Is there evidence of racial disparity in police use of deadly force? Analyses of officer-involved fatal shootings in 2015–2016. *Social Psychological Personality Science*, 10, 586–595.
- Fyfe, J. J. (1982). Blind justice: Police shootings in Memphis. *J. Crim. L. Criminol.*, 73, 707–722.
- Geller, W. A., & Karales, K. J. (1981). Shootings of and by Chicago police: Uncommon crises—Part I: Shootings by Chicago police. *J. Crim. L. Criminol.*, 72, 1813–1866.
- 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.
- Hehman, E., Flake, J. K., & Calanchini, J. (2018). Disproportionate use of lethal force in policing is associated with regional racial biases of residents. *Soc. Psychol. Personal. Sci.*, 9(4), 393–401.
- Inn, A., Wheeler, A. C., & Sparling, C. L. (1977). The effects of suspect race and situation hazard on police officer shooting behavior. *J. Appl. Soc. Psychol.*, 7(1), 27–37.
- Jacobs, D. (1998). The determinants of deadly force: A structural analysis of police violence. *Am. J. Sociol.*, 103(4), 837–862.
- 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.
- Menard, S. (2002). *Applied logistic regression analysis* (Vol. 106). Sage.
- Nix, J., Campbell, B. A., Byers, E. H., & Alpert, G. P. (2017). A bird's eye view of civilians killed by police in 2015: Further evidence of implicit bias. *Criminology & Public Policy*, 16(1), 309–340.
- Ross, C. T. (2015). A multi-level bayesian analysis of racial bias in police shootings at the county-level in the United States, 2011–2014. *PLOS One*, 10(11), e0141854.
- Schimmack, U. (2019). *Does PNAS article show there is no racial bias in police shootings?* Retrieved from <https://replicationindex.com/2019/09/28/does-pnas-article-show-there-is-no-racial-bias-in-police-shootings/> (Online; posted 09/28/2019)
- Scott, K., Ma, D. S., Sadler, M. S., & Correll, J. (2017). A social scientific approach toward understanding racial disparities in police shooting: Data from the Department of Justice (1980–2000). *J. Soc. Issues*, 73(4), 701–722.
- Takagi, P. (1974). A garrison state in 'democratic' society. *Crime Soc. Justice*, 27–33.
- te Grotenhuis, M., Pelzer, B., Eisinga, R., Nieuwenhuis, R., Schmidt-Catran, A., & Konig, R. (2017). When size matters: Advantages of weighted effect coding in observational studies. *International Journal of Public Health*, 62(1), 163–167.