Title: High-frequency location data show that race affects citations and fines for speeding

Authors: Pradhi Aggarwal[1], Alec Brandon[2*], Ariel Goldszmidt[3], Justin Holz[4], John A. List[3], Ian Muir[5], Gregory Sun[6], Thomas Yu[7]

Affiliations:

[1]Massachusetts Institute of Technology; Cambridge, Massachusetts, USA.

[2]Johns Hopkins University; Baltimore, Maryland, USA.

[3]University of Chicago; Chicago, Illinois, USA.

[4]University of Michigan; Ann Arbor, Michigan, USA.

[5]Wal-Mart Inc.; Bentonville, Arkansas, USA.

[6]Washington University; St. Louis, Missouri, USA.

[7]Yale University; New Haven, Connecticut, USA.

*Corresponding author. Email: alec.brandon@jhu.edu

Abstract: Prior research on racial profiling finds that in encounters with law enforcement, minorities are punished more severely than white civilians. Less is known about the causes of these encounters and their implications for our understanding of racial profiling. Using high-frequency location data of rideshare drivers in Florida (*N*=222,838), we estimate the effect of driver race on citations and fines for speeding using 19.3 million location pings. Compared to a white driver traveling the same speed, we find that racial/ethnic minority drivers are 24 to 33 percent more likely to be cited for speeding and pay 23 to 34 percent more money in fines. We find no evidence that accident and re-offense rates explain these estimates, suggesting that underlying our results is an animus against minorities.

Racial profiling is one of the most significant social issues in the United States. To better understand the extent of this issue and motivate policy interventions to mitigate it, social scientists have extensively examined administrative data collected by law enforcement agencies, such as records of police encounters with motorists and pedestrians, responses to 911 calls, and criminal trials. Broadly speaking, this research finds that, relative to the white civilians who appear in the data, police officers and judges are more likely to search, fine, use force, charge, detain, and incarcerate civilians in the dataset who belong to a racial or ethnic minority group (*1–17*).

A basic concern with this research is that, because race and ethnicity likely affect whether a civilian is cited in the administrative data, the results are unlikely to capture the full extent of racial profiling (*18–23*). The direction in which this would bias prior findings depends on why race or ethnicity affects the likelihood of appearing in the data. One possibility is that civilians who belong to a racial or ethnic minority group are more likely to be cited than their white counterparts for the same violation of the law, potentially because they are targeted by police, or their communities are overpoliced. Under this possibility, analyses of administrative records would understate the extent of racial profiling. Another possibility is that minority civilians are more likely to be cited because they break the law more frequently, which would cause analyses of the administrative data to overstate the extent of racial profiling.

Without additional data that allow social scientists to control for possibilities like these, it is hard to make progress on policy interventions to mitigate racial profiling. A subset of the prior research suggests that interventions should target "taste-based discrimination" (*24*) because underlying the evidence of racial profiling is an animus against racial and ethnic minority groups, not "statistical discrimination" (*25,26*) where police and judges use civilian race and ethnicity as predictors that help accomplish objectives like deterrence. However, it is possible this suggestion is only partially true and that an analysis of the full extent of racial profiling would find evidence for targeting statistical discrimination too.

In this article, we examine the full extent of racial profiling in the enforcement of speed limits. Prior research on this question is emblematic of the broader literature on racial profiling. Among motorists cited in the administrative data, minority drivers receive harsher punishments for speeding (*5,7,13–15*) and these harsher punishments were attributed to taste-based, not statistical, discrimination

because they cannot be explained by an objective like deterring future citations for speeding (*14*). There is also evidence that minority drivers are more likely to appear in these data, although why they are overrepresented is multivariate. On the one hand, there is evidence that police officers target minority drivers in traffic stops (*13*) and that minority communities are overpoliced (*27,28*). On the other hand, there are also data collected by researchers (*29*) and automated speed cameras (*30*) that suggest drivers who belong to a racial or ethnic minority group are more likely to exceed the speed limit. One way to better understand the full extent of racial profiling in this context would be to analyze administrative records on speeding violations that are linked to data on whether a motorist was driving, their driving speed relative to the limit, their race and ethnicity, and other potentially important controls like driving location.

We construct a new dataset of this form by linking (1) high-frequency location or HFL data (i.e., global positioning system or GPS pings communicated from a driver's smartphone) of motorists with (2) official government administrative records on their speeding violations and (3) inference of their race and ethnicity. To the best of our knowledge, this is the first study to construct such a dataset (*31*). Using this data, we conduct analyses to answer the following research questions (RQs).

RQ1: Compared to white drivers, are minority drivers more likely to be cited and pay more in fines for speeding?

RQ2: Which mechanism, taste-based or statistical discrimination, underlies our answer to RQ1?

Method

We answer these questions using data from a collaboration with the ride-share platform, Lyft. This study's sample included Lyft drivers (N=222,838) operating in the state of Florida. Much like its competitor Uber, Lyft runs a platform that matches customers with drivers offering rides. To offer rides, drivers log into an application on their smartphone that every ten seconds communicates pings of a driver's precise location and driving speed to Lyft's platform.

Measures

Driver speed/re-offense of speeding/reckless driving: The HFL pings that Lyft receives denote the time and location of a driver using their smartphone's clock and GPS. We calculate average driving

speed between pings by dividing the distance and time between each ping. After merging the location of the pings onto Florida records of the location of road segments and their attributes, such as speed limit, we can infer whether and the extent to which a driver exceeds the speed limit. Following the relationship between speeding and fines in Florida, we group speeding into the following discrete categories: 0 to 9, 10 to 14, 15 to 19, 20 to 29, and 30+ miles per hour (mph) over the speed limit.

Citations/fines/accidents: In partnership with Lyft, we obtained access to the pings of the population of Lyft drivers operating in Florida between August 2017 and August 2020. We link each ping to Florida's government records of speeding violations and accidents with information on timing and driver's license number. We infer when an accident occurred or when a driver was stopped by police for a speeding citation by searching for the largest reduction in driving speed within three minutes of the approximate time in the Florida records.

Race/ethnicity: To infer driver race/ethnicity, we fit a model that relates the self-reported race/ethnicity of the 45.7 percent of drivers who are registered voters in Florida to a picture each driver submits when they apply to drive on Lyft's platform. This fitted model is then used to infer the race/ethnicity from the picture submitted by the 54.3 percent of drivers who are not registered voters. Our main analysis assigns drivers to one of two racial and ethnic categories: white or minority, where minority includes Asian and Pacific Islander, Black, or Hispanic (*32*). For more on the construction of our analysis sample, we present a longer discussion and summary statistics in SM section 1.1-6.

To determine whether there is evidence of racial profiling, we would ideally compare white and minority drivers while controlling for all other factors that are related to driver race and also affect speeding citations and fines. Under the first strategy, we select the controls that are likely relevant. We refer to this strategy as the "FE Model" because we primarily select fixed effects to control for the non-race factors that could influence the receipt of citations and fines. Second, a machine learning algorithm selects the factors to control. We refer to this strategy as the "DML Model" because it relies on a technique called double machine learning (*33*). See SM 2.1 for more information about the variables used in the FE Model and see SM 2.2 for more on the DML Model's selection of controls.

Taste-based versus statistical discrimination

To understand whether the racial profiling of minorities is attributable to statistical or taste-based discrimination by police, we adopt the following strategy. We test for statistical discrimination

by examining whether, compared to white drivers, speeding citations for minority drivers are more predictive of other traffic violations. If we find evidence that minority status is predictive, then we would attribute racial profiling to statistical discrimination because it would suggest police profile to deter traffic violations that happen to be more prevalent amongst minority drivers than white motorists. Otherwise, we would attribute the evidence to taste-based discrimination.

We consider two types of traffic violations with which a citation for speeding is likely related: reckless driving (using accidents as a proxy) and re-offenses (of speeding). With accidents, we test for whether minority status is predictive by regressing whether an accident occurs on the interaction of minority status, speeding, and the controls used in the FE Model described above. With re-offenses, we test for whether minority status is predictive by regressing the portion of time in a week spent driving more than 10 miles per hour over the speed limit on the interaction between minority status and weeks since a citation for speeding, driver fixed effects, and weeks since a citation fixed effects (*34*). If both tests find that minority status is a statistically significant predictor, then we would attribute an effect of minority race/ethnicity to statistical discrimination by police; in contrast, if both tests find no such statistical significance, then we would attribute an effect of minority status to taste-based discrimination.

Results

Our empirical analysis of racial profiling controls for location, time, features of the road, car, driver age and gender, the phase of a trip the driver was completing (i.e., waiting to be paired with a passenger, picking up a passenger, or dropping off a passenger), and the extent to which the driver was exceeding the speed limit (see SM 2.1-2.2 for details).

RQ1: Race, Citations, and Fines

We find that compared to white drivers, minority drivers are significantly more likely to be cited for speeding and they pay significantly more in fines. Importantly, this comparison controls for whether and the extent to which a driver was speeding, where they were driving, features of the car being driven, and other potentially relevant covariates. Figure 1 presents these findings by plotting the rates of citations and fines for white and minority drivers, as well as the estimated effect of driver race on citations and fines for speeding. Across the two outcomes and models in the top row of Fig. 1, we see that the average rate of citations and cost of fines is low. For every 10,000 hours of driving, which is

about a half decade of working 40 hours per week, drivers have an 11 to 15 percent chance of receiving a citation for speeding and the average cost of these citations is $15 to $20. Yet despite how rare speeding citations are, differences between white and minority drivers appear across both outcomes and both models.

The bottom row of Fig. 1 plots the average effect of minority status on citation rates and fines. Starting from the left, the FE and DML models respectively find that for every 10,000 hours of driving, minority drivers are 3.6 and 2.7 percentage points more likely to be cited for speeding than white drivers. Moving to the right, we report a similar relationship between minority status and fines, with the FE and DML models respectively finding that the average effect of minority status on speeding fines per 10,000 of driving is $5.20 and $3.58. Relative to how rare citations and fines are, these effects are quite large. In percentage terms, minority drivers are 24 to 33 percent more likely to be cited for speeding and pay 23 to 34 percent more in fines. Furthermore, the four estimated effects on the bottom two panels of Fig. 1 are statistically significantly different from zero at the conventional level of five percent. SM 3.1-3.2 present these results in table form. The results from a series of robustness and heterogeneity analyses can be found in SM 4.1-4.6, these include the effect of minority status at discrete levels of speeding and the effect of each race/ethnicity.

RQ2: Mechanisms—Taste-based versus statistical discrimination

Having established that minority drivers are more likely to be cited and fined for speeding than white drivers, we next consider mechanisms (i.e., taste-based versus statistical discrimination). Our analysis finds that, contrary to the statistical discrimination mechanism, police do not racially profile in their enforcement of speeding tickets to deter reckless driving (accidents) or reoffending speed limits. In particular, we find that white and minority drivers have statistically indistinguishable accident rates and rates of reoffending, suggesting that statistical discrimination does not underlie the evidence of racial profiling. As a result, we attribute the effect of minority status on citations and fines to taste-based discrimination by police officers. Figure 2 presents the results of an analysis that considers whether police officers are more likely to penalize minority drivers for speeding because their speeding is more predictive of accidents than for white drivers. The effect of minority status on accident rates is reported for several discrete levels of driving speed exceeding the speed limit with brackets around each estimated effect illustrating the 95 percent confidence interval. When travelling 0 to 9 mph over the

speed limit, minority drivers are marginally more likely to be cited for an accident than white drivers; at nearly every higher level of speeding, the direction reverses, and white drivers are more likely to be cited for an accident. Furthermore, none of these estimates are statistically distinguishable from zero at the conventional level of 5 percent significance.

Figure 3 presents the results of a complementary analysis that considers whether police officers are more likely to penalize minority drivers because their speeding is more predictive of reoffending. This figure shows that being a minority does not predict speeding around a citation, as the difference has an ambiguous sign in the weeks before and after a speeding ticket is received. Furthermore, the differences are statistically indistinguishable from zero at the conventional level of 5 percent significance. Regression results can be found in SM 3.3-3.4.

Discussion

Using data on active Lyft drivers in Florida, we estimate the effect of racial profiling on citations and fines for speeding. Across two empirical strategies that control for speeding, location, and a battery of other covariates, we find that racial/ethnic minority drivers are significantly more likely to be cited for speeding and also pay significantly more money in fines. Our complementary analyses of re-offense and accident rates fail to find evidence that these effects can be attributed to differences in speeding behavior or traffic violations across white versus minority drivers (i.e., statistical discrimination); therefore, findings suggest that police racial profiling of drivers is due to police having animus or prejudice against minority drivers (i.e., taste-based discrimination).

These results suggest concrete prescriptions for policymakers and business leaders interested in mitigating racial profiling. For policymakers, our findings suggest that, relative to enforcement by police officers, appropriately located automated technologies such as speeding cameras could help reduce selective enforcement of traffic violations. For business leaders, our findings indicate that race-blind approaches to disincentivizing risky behavior may not be so blind. For example, car insurance rates typically increase when drivers are cited for speeding (*35*), but our findings indicate that such citations are not blind to driver race. Accounting for race in the relationship between citations and insurance rates could help diminish the impact of racial profiling.

More generally, we present a methodological advance. We view our analysis as a proof of concept for a new approach to measuring the effects of racial profiling. A basic concern with prior

research is that it can only evaluate racial profiling after an encounter with police has begun; however, failing to control for the previous events that lead to these encounters (e.g., speeding) can bias evaluations for or against finding evidence of racial profiling. Our analysis demonstrates that these events can be controlled in the context of enforcement of speed limits. With sufficiently large datasets of motorists or pedestrians, it is likely possible to extend our approach to the analysis of outcomes that follow traffic and pedestrian stops, such as use of force, detainment, and incarceration. While citations for speeding are too rare in our data to investigate outcomes that follow, progress would likely be possible with larger samples of HFL. A similar analysis of racial profiling in the setting of stop-and-frisk would likely be possible with HFL data that could link pedestrians to their encounters with police.

We also conjecture that, with a more representative sample of HFL data, applications of our approach would be superior to existing alternatives. The "veil of darkness" test (*13,19*), which compares the racial composition of citations before and after the sun sets is useful for determining whether there is more racial profiling during evening hours but cannot address profiling at other times of the day. Similarly, a more recent alternative compares the racial composition of citations issued by police to those issued by automated cameras (*30,36*); however, this approach is also limited because it only evaluates racial profiling in the locations where cameras are installed.

Limitations

This article has several important limitations. First, we study a very specific sample of motorists that is almost certainly not representative of drivers in Florida or U.S. For example, to drive on Lyft's platform, motorists must pass a background check that excludes applicants convicted of violent offenses, sexual offenses, and driving under the influence of drugs or alcohol. Lyft drivers are incentived to avoid traffic violations and these incentives work. Compared to the general population of motorists, our sample is less prone to speed, especially more than 10 mph over the limit (*27*). As a result, our analysis examines only 1,423 citations for speeding, which limited the scope of analyses we could conduct. With more citations, researchers could build on our analyses to learn more about the nature of the taste-based discrimination that we find underlies racial profiling. Of particular interest would be analyses that precisely estimate the effect of driver race/ethnicity on speeding citations as a function of ambient light, neighborhood demographics, and motor vehicle characteristics.

Second, our sample of Lyft drivers is observed in Florida, during 2017 to 2020. In terms of racial and ethnic demographics, Florida is similar to the rest of the country (*37*) but is more politically conservative, and Florida's police are less likely to search a car during a stop; those searches largely target non-Hispanic drivers (*13*). While we only study a specific timeframe, the SM shows that omitting data after COVID-19 lockdowns and protests over George Floyd's killing did not impact results.

Third, we still rely on some administrative data from law enforcement that may mask important forms of racial profiling. Most significant is the possibility that police officers use civilian race as a pretext to stop civilians and not cite these stops in their records. However, we conjecture that progress on unrecorded stops would require combining data like ours with data reported by drivers or HFL data on police officers (*28*). A related limitation of our HFL data is that it is proprietary and can only be accessed for use with permission from Lyft.

References and Notes

1. J. Knowles, N. Persico, P. Todd, Racial Bias in Motor Vehicle Searches: Theory and Evidence. *Journal of Political Economy* **109**, 203-229 (2001). doi:10.1086/318603
2. S. Anwar, H. Fang, An Alternative Test of Racial Prejudice in Motor Vehicle Searches: Theory and Evidence. *American Economic Review* **96**, 127-151 (2006). doi:10.1257/000282806776157579
3. A. Gelman, J. Fagan, A. Kiss, 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**, 813-823 (2007). doi:10.1198/016214506000001040
4. K. Antonovics, B. Knight, A New Look at Racial Profiling: Evidence from the Boston Police Department. *Review of Economics and Statistics* **91**, 163-177 (2009). doi:10.1162/rest.91.1.163
5. M. Makowsky, T. Stratman, Political Economy at Any Speed: What Determines Traffic Citations? *American Economic Review* **99**, 509-527 (2009). doi:10.1257/aer.99.1.509
6. D. Abrams, M. Bertrand, S. Mullainathan, Do Judges Vary in Their Treatment of Race? *Journal of Legal Studies* **41**, 347-383 (2012). doi:10.1086/666006
7. N. Anbarci, J. Lee, Detecting racial bias in speed discounting: Evidence from speeding tickets in Boston. *International Review of Law and Economics* **38**, 11-24 (2014). doi:10.1016/j.irle.2014.02.002
8. M. Makowsky, T. Stratman, More Tickets, Fewer Accidents: How Cash-Strapped Towns Make for Safer Roads. *Journal of Law and Economics* **54**, 863-888 (2011). doi:10.1086/659260
9. A. Gupta, C. Hansman, E. Frenchman, The Heavy Costs of High Bail: Evidence from Judge Randomization. *Journal of Legal Studies* **45**, 471-505 (2016). doi:10.1086/688907
10. D. Arnold, W. Dobbie, C. Wang, Racial Bias in Bail Decisions. *Quarterly Journal of Economics* **133**, 1885-1932 (2018). doi:10.1093/qje/qjy012

11. R. Fryer, An Empirical Analysis of Racial Differences in Police Use of Force. *Journal of Political Economy* **127**, 1210-1261 (2019). doi:10.1086/701423
12. J. West, Racial Bias in Police Investigations. Working Paper (2018).
13. E. Pierson, C. Simoiu, J. Overgoor, et al., A large-scale analysis of racial disparities in police stops across the United States. *Nature Human Behaviour* **4**, 736-745 (2020). doi:10.1038/s41562-020-0858-1
14. F. Goncalves, S. Mello, A Few Bad Apples? Racial Bias in Policing. *American Economic Review* **111**, 1406-1441 (2021). doi:10.1257/aer.20181607
15. S. Anwar, P. Bayer, R. Hjalmarsson, M. Mizel, Racial Disparities in Misdemeanor Speeding Convictions. *RAND Corporation Report* (2021).
16. M. Hoekstra, C. Sloan, Does Race Matter for Police Use of Force? Evidence from 911 Calls. *American Economic Review* **112**, 827-860 (2022). doi:10.1257/aer.20201292
17. D. Arnold, W. Dobbie, P. Hull, Measuring Racial Discrimination in Bail Decisions. *American Economic Review* **112**, 2992-3038 (2022). doi:10.1257/aer.20201653
18. I. Ayres, J. Waldfogel, A Market Test for Race Discrimination in Bail Settings. *Stanford Law Review* **46**, 987-1047 (1994). doi:10.2307/1229062
19. J. Grogger, G. Ridgeway, Testing for Racial Profiling in Traffic Stops From Behind a Veil of Darkness. *Journal of the American Statistical Association* **101**, 878-887 (2006). doi:10.1198/016214506000000168
20. C. Simoiu, S. Corbett-Davies, S. Goel, The problem of infra-marginality in outcome tests for discrimination. *Annals of Applied Statistics* **11**, 1193-1216 (2017). doi:10.1214/17-AOAS1058
21. S. Duflauf, J. Heckman, An Empirical Analysis of Racial Differences in Police Use of Force: A Comment. *Journal of Political Economy* **128**, 3679-4009 (2020). doi:10.1086/710976
22. I. Canay, M. Mogstad, J. Mountjoy, On the Use of Outcome Tests for Detecting Bias in Decision Making. *Review of Economic Studies* **91**, 2135-2167 (2024). doi:10.1093/restud/rdad082
23. D. Knox, W. Lowe, J. Mummolo, Administrative Records Mask Racially Biased Policing. *American Political Science Review* **114**, 619-637. doi:10.1017/S0003055420000039
24. G. Becker, The Economics of Discrimination, University of Chicago Press (1957).
25. K. Arrow, The Theory of Discrimination. *Discrimination in Labor Markets*, O. Ashenfelter and A. Rees, eds., Princeton, NJ: Princeton University Press (1972).
26. E. Phelps, The Statistical Theory of Racism and Sexism. *American Economic Review* **62**, 659-661 (1972).
27. W. Cai, J. Gaebler, J. Kaashoek, L. Pinals, S. Madden, S. Goel, Measuring racial and ethnic disparities in traffic enforcement with large-scale telematics data. *PNAS Nexus* **1**, 1-8 (2022). doi:10.1093/pnasnexus/pgac144
28. M. Chen, K. Christensen, E. John, E. Owens, Y. Zhao, Smartphone Data Reveal Neighborhood-Level Racial Disparities in Police Presence. *Review of Economics and Statistics* (Forthcoming). doi:10.1162/rest_a_01370
29. R. Tillyer, R. Engel, Racial differences in speeding patterns: Exploring the differential offending hypothesis. *Journal of Criminal Justice* **40**, 285-295 (2012). doi:10.1016/j.jcrimjus.2012.04.001
30. W. Xu, M. Smart, N. Tilahun, et al., The racial composition of road users, traffic citations, and police stops. *Proceedings of the National Academy of Arts and Sciences* **121**, 1-6 (2024). doi:10.1073/pnas.2402547121

31. Earlier research with HFL data did not link the data to individual-level data on race/ethnicity and speeding violations (*27,30*).
32. Registration forms in Florida have voters indicate whether they are White, Not Hispanic or Hispanic, so White Hispanics are classified as Hispanic in our data.
33. V. Chernozhukov, D. Chetverikov, M. Demirer, E. Duflo, C. Hansen, W. Newey, J. Robins, Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal* **21**, C1-C68 (2018). doi:10.1111/ectj.12097
34. Our re-offenses test depends on whether minority status is predictive of speeding both before and after a citation because, in isolation, indistinguishable post-citation differences in speeding could be a consequence of racial profiling more powerfully deterring the speeding of minority drivers. However, if these differences are indistinguishable both before and after a citation, then it is not straightforward to reconcile with a statistically motivated deterrence policing strategy.
35. A. Carrns, What That Speeding Ticket May Really Cost You. New York Times (2012). url:https://archive.nytimes.com/bucks.blogs.nytimes.com/2012/05/17/what-that-speeding-ticket-may-really-cost-you/
36. S. Quintanar, Man Vs. Machine: An Investigation of Speeding Ticket Disparities Based on Gender and Race. *Journal of Applied Economics* **20**, 1-28 (2017). doi:10.1016/S1514-0326(17)30001-6
37. A. Van Dam, What state best represents America? Washington Post, (2024). url:https://www.washingtonpost.com/business/2024/05/10/most-representative-most-unique-places-america/

Supplementary materials references

38. S. Lundberg, S. Lee, A Unified Approach to Interpreting Model Predictions. *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 4768-4777 (2017).
39. S. Lundberg, G. Erion, S. Lee, Consistent Individualized Feature Attribution for Tree Ensembles. *Arxiv Working Paper*, (2018). doi:10.48550/arXiv.1802.03888

Acknowledgments: Bocar Ba, Roland Fryer, Jeffery Grogger, and Magne Mogstad provided helpful comments and suggestions. Ashley Adams, Jon Alferness, Klaus Ashorn, Stephanie Carpenter, Nicholas Chamandy, Peter Day, Brian Emeott, Keith Falcon, Heather Foster, Logan Green, Tommy Hayes, Srishti Kapur, Chris Martin, Brian McClure, Shadawn Reddick-Smith, Betsey Stevenson, Amber Wang, and John Zimmer provided invaluable assistance in advancing this study.

Funding: None.

Author contributions

Conceptualization: AB, JH, JAL, TY

Methodology: AB, PA, AG, JH, JAL, GS, TY

Investigation: JAL, IM, TY

Visualization: PA, TY

Funding acquisition: None

Project administration: AB, JH, IM

Supervision: JAL, IM

Writing – original draft: AB, JH, TY

Writing – review & editing: AB, JAL, TY

Competing interests: This study was conducted while Pradhi Aggarwal, Ariel Goldszmidt, John List, Ian Muir, Gregory Sun, and Thomas Yu were employed by Lyft. None are currently employed by Lyft.

Data and materials availability: All code and non-proprietary data used to conduct our analysis are stored on Dryad and accessible at the following hyperlink: https://doi.org/10.5061/dryad.4f4qrfjnk. The non-proprietary data are stored on Dryad in a redacted form that removes personally identifying information. The proprietary data that are not stored on Dryad are the HFL data from Lyft (i.e., the smartphone pings from active Lyft drivers) and the voter data, which was obtained from a third-party vendor. A full replication of the merges that created the dataset used to conduct our analysis is not guaranteed because Lyft does not store its HFL data indefinitely. The dataset used to conduct our analysis remains on Lyft's servers and we have requested that Lyft retain the data for at least five years. Researchers seeking replication with this dataset should submit an emailed request to bemeott@lyft.com and kfalcon@lyft.com clearly stating the analyses they would like to reproduce, the names and organizational affiliations of the researchers, and a proposed data use agreement. Lyft will review all such requests but will grant access at its sole discretion. If the aforementioned email addresses become disabled, then please contact A. Brandon, the corresponding author.

Supplementary Materials

Materials and Methods

Figs. S1 to S5

Tables S1 to S13

References (*37*–*38*)

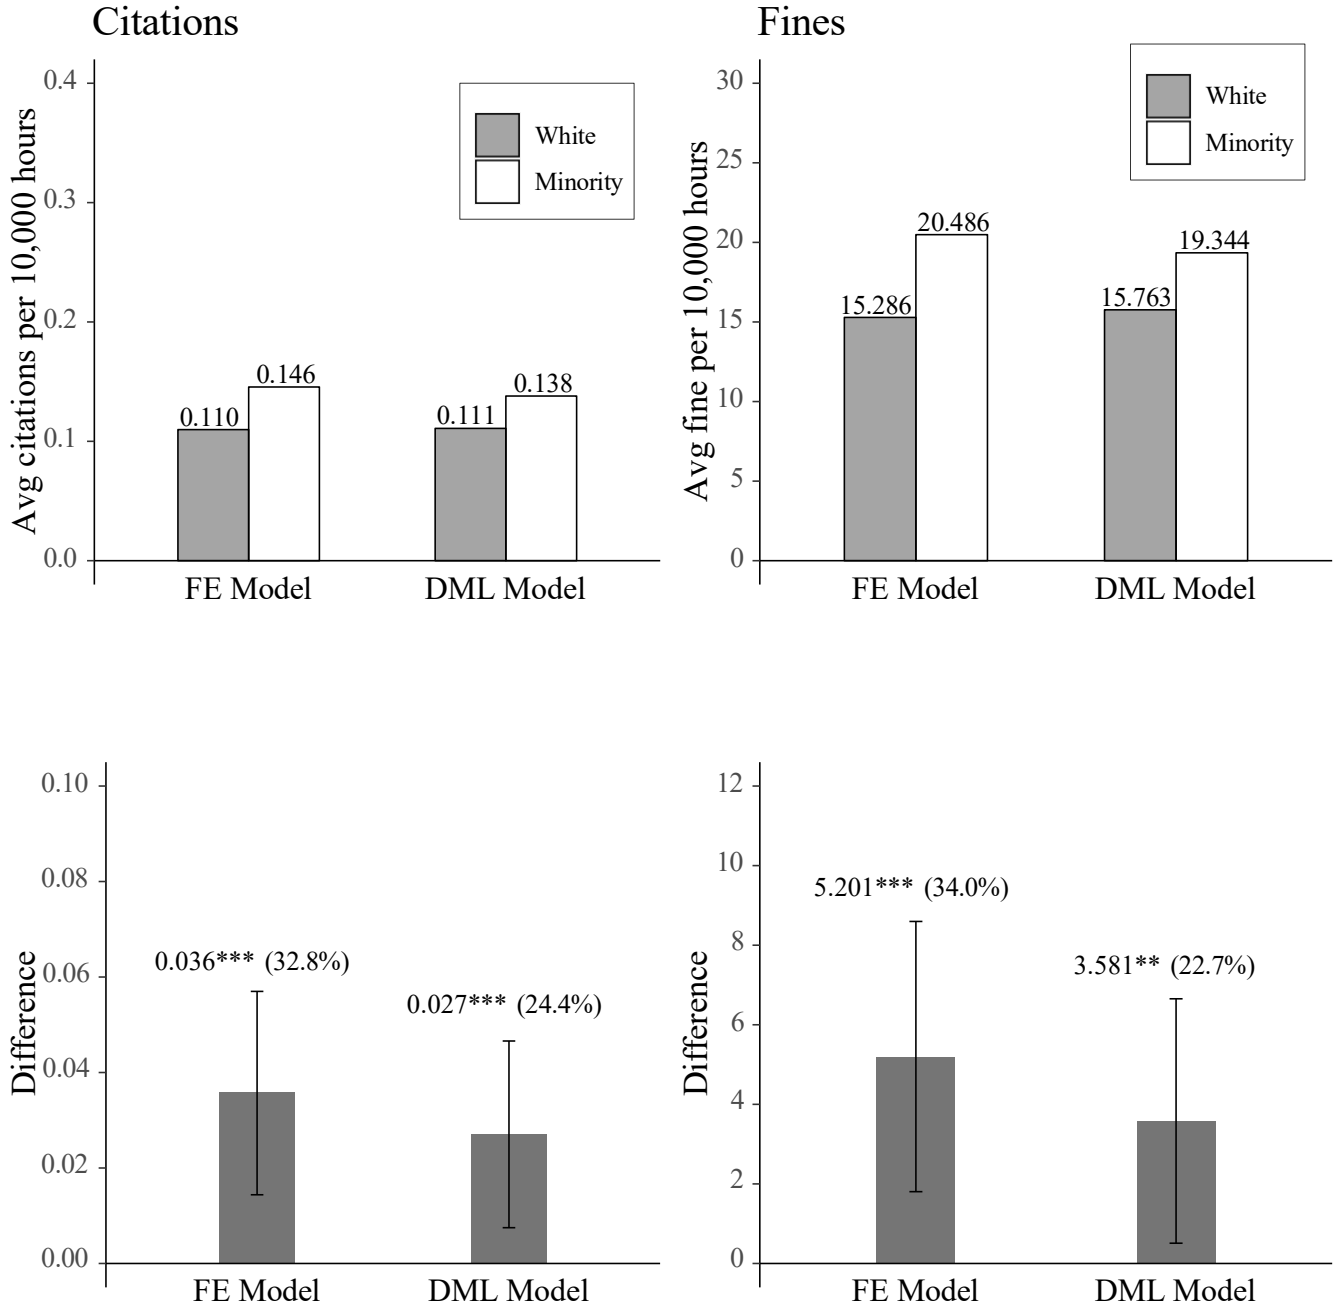


Fig. 1. Effect of Driver Race on Citations and Fines for Speeding. In the top row of this figure the rate of citations and fines for white and minority drivers is plotted. In each column of the top row, the left pair of estimates are adjusted for the variation explained by the control variables in the FE model, while the right pair of estimates account for variation explained by controls in the DML model. In the bottom row of this figure, the difference between white and minority citations rates and fines are plotted. For each outcome, the left estimate controls for the covariates in the FE model and the right estimate controls for the covariates selected by the DML model. 95 percent confidence intervals constructed with robust standard errors clustered by driver are included on the estimated differences. In parentheses the estimated difference is normalized by the rate of citation and fines for white drivers. *** p-value < 0.01, ** p-value < 0.05, * p-value < 0.10.

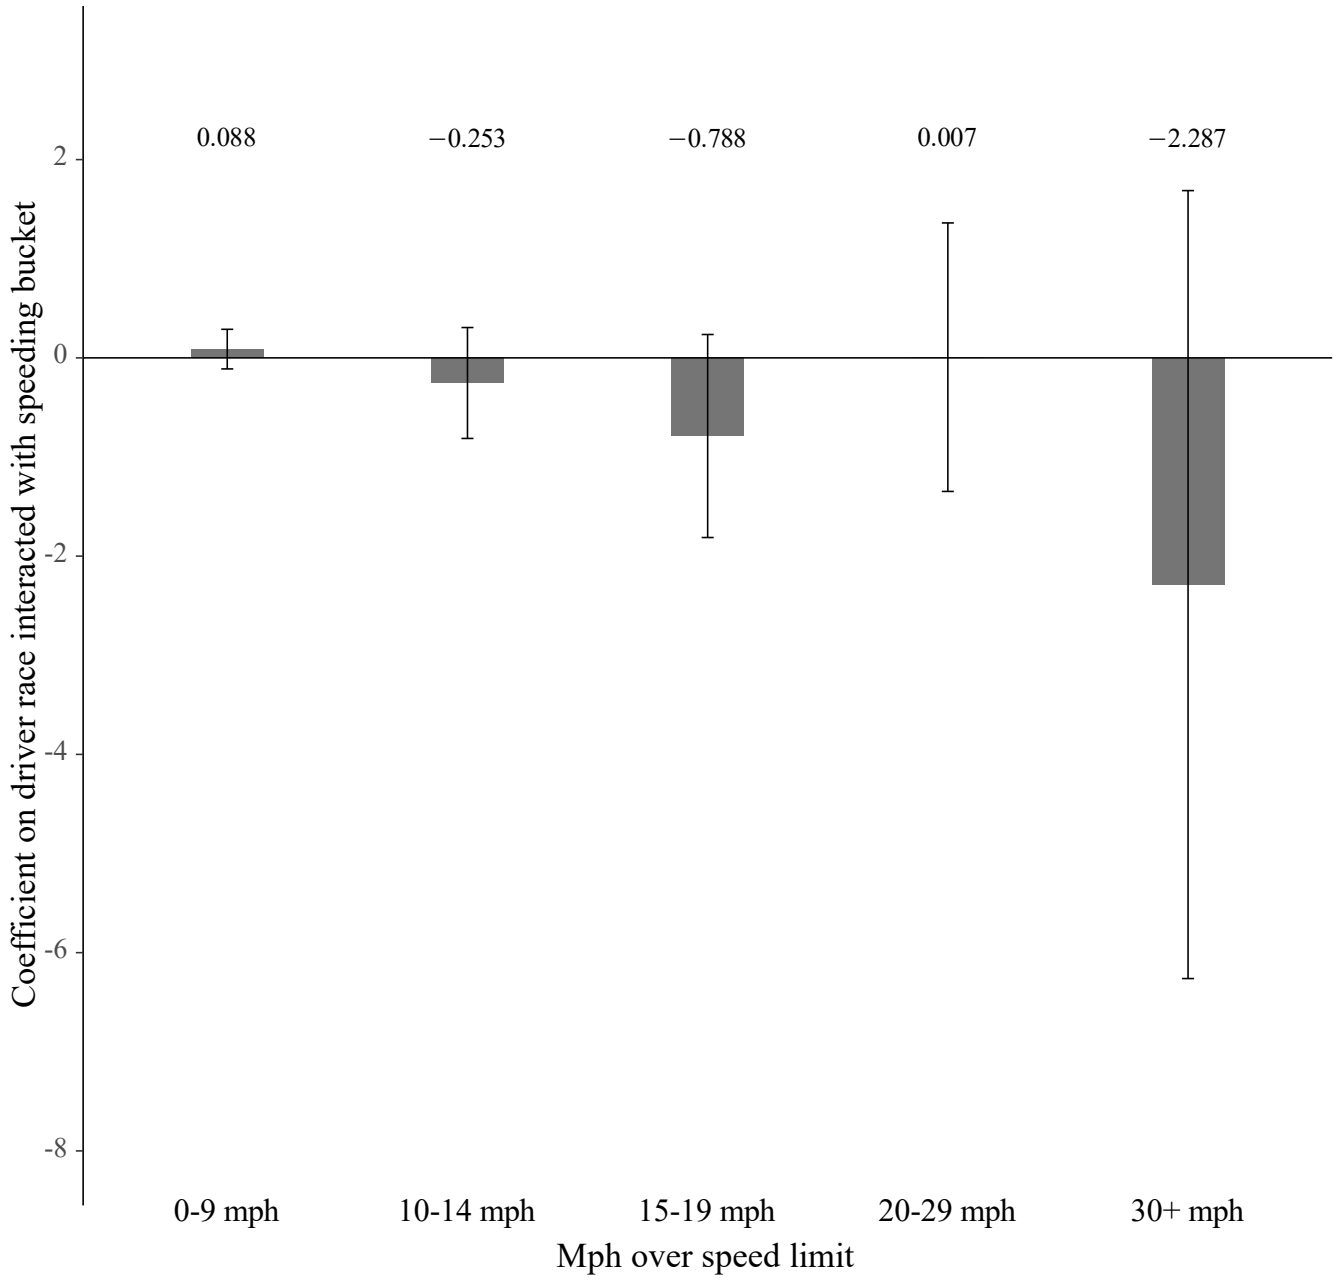
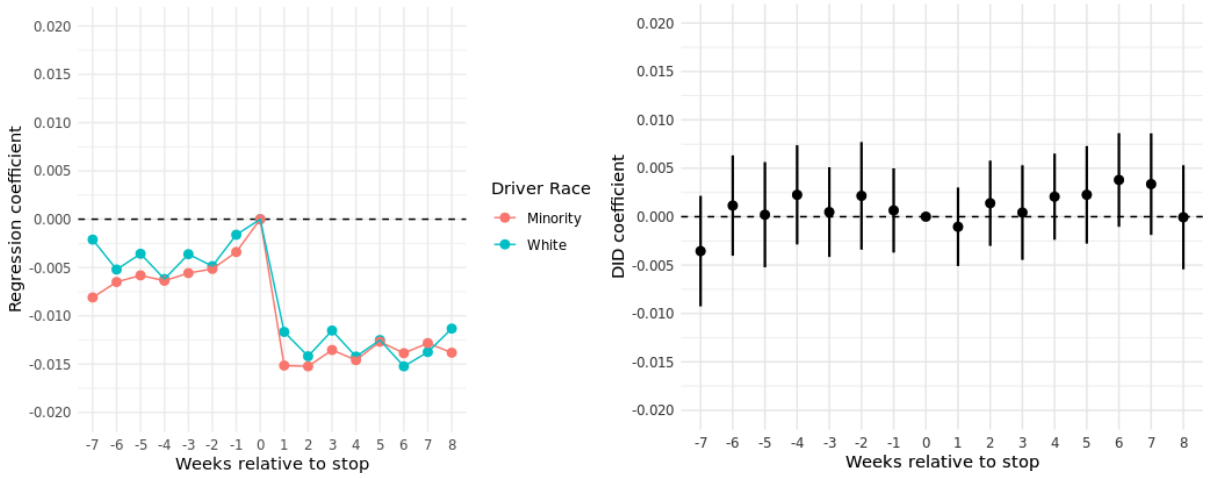


Fig. 2. Effect of Driver Race on Relationship Between Driving Speed and Accidents. This figure plots the difference between white and minority driver accident rates by level of speeding. The estimated differences control for the covariates in the FE model. Brackets indicate the 95 percent confidence interval, which are constructed with standard errors clustered by driver.

Figure 4:
Effect of Citation on Reoffenses



(A) Re-offense by Race

(B) Racial Differences in Reoffenses

Fig. 3. Effect of Speeding Citations on Reoffending. Panel A plots the proportion of driving time over the speed limit in the weeks before and after a white and minority driver is cited for a speeding ticket. Estimates are relative to time spent speeding during the week a speeding ticket was received and control for driver fixed effects. Panel B plots the estimated differences with 95 percent confidence intervals constructed with robust standard errors clustered by driver.

Supplementary materials for "High-frequency location data show that race affects citations and fines for speeding"

Pradhi Aggarwal,[1] Alec Brandon,[2*] Ariel Goldszmidt,[3] Justin Holz,[4]
John A. List,[3] Ian Muir,[5] Gregory Sun[6],Thomas Yu[7]

[1]Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA
[2]Johns Hopkins University, Baltimore, Maryland 21202, USA
[3]University of Chicago, Chicago, Illinois 60637, USA
[4]University of Michigan, Ann Arbor, Michigan 48109, USA
[5]Wal-Mart Inc.; Bentonville, Arkansas, USA 72716, USA
[6]Washington University, St. Louis, Missouri 63105, USA
[7]Yale University, New Haven, Connecticut 06520, USA

[*]To whom correspondence should be addressed; E-mail: alec.brandon@jhu.edu.

1 Analysis Dataset Construction

This section describes the data sources and how they were combined to create our main analysis dataset.

1.1 High-Frequency Lyft Data

We obtain the universe of pings sent from Lyft drivers operating in Florida between August 30th 2017 to August 25th 2020. Calculating the distance between pings and the time between pings allows us to measure driving speed. Merging the location of these pings to Florida road segment and speed limit records allows us to infer the prevailing speed limit for each ping.

Supplementary Appendix Table S1 presents summary statistics on the prevalence off speeding for the Lyft drivers in our analysis sample by race and ethnicity.

The Lyft data also allows us to observe the driver's license number. These numbers are likely accurate because drivers have to pass a Department of Motor Vehicles background check to operate on the platform.

1.2 Citation Data

We examine the the universe of citations written by the Florida State Highway Patrol from August 30th 2017 to August 25th, 2020. The first column of Supplementary Appendix Table S2 summarizes the demographic variables of cited drivers in this data. To prepare these data for a merge with the high-frequency Lyft data, we drop citations that do not meet the following criteria:

1. The stop is due to a speeding citation.
2. The observation is not missing a driver's license number.
3. The cited driver's license number matches the driver's license number of a driver who was active in Florida between August 30th 2017 to August 25th 2020.
4. The time of the citation occurs within 5 minutes of the motorist actively driving on Lyft's platform.

This process results in 1,423 citations.

1.3 Inferring the Cited High-Frequency Location Ping

To determine the speeding behavior that led to each observed citation, we must link each citation to a specific high-frequency location ping. To link these citations to the specific ping that

triggered the citation we consider all location pings that occur in the 3 minutes leading up to the citation. We match each citation to a single ping using the following criteria in order of preference:

1. The ping where the road speed limit matches the police reported speed limit
2. The ping with the fastest speed
3. The ping that occurs closest to the estimated stop time

Supplementary Appendix Table S1 presents summary statistics on the rate of speeding citations and magnitude of fines for speeding for the Lyft drivers in our analysis sample by race and ethnicity.

1.4 Accident Data

The information on crash reports during our sample period comes from the Florida Department of Transportation's (FDOT) administrative records of known crashes. FDOT collects this information during a police response or investigation. The data also contains the date and time of the incident, information on injuries to individuals involved in the crash, and an estimated dollar value of property damage.

FDOT identifies the drivers involved in the accident by their driver's license number. During the sample period, we observe 8,243 accidents that we are able to match to Lyft drivers while they are online on the Lyft platform. Note that not all of these drivers receive a speeding citation in our dataset. We use the same process as described in the previous section to match accidents to Lyft location pings.

1.5 Race Prediction Model

We build a race prediction model with four components. First, we train a convolutional neural network (CNN) to predict race from profile pictures of Lyft users. These profile pictures are obtained from the 5,531 users in Florida who responded to Lyft's 2019 Economic Impact Report (EIR) survey and reported their race. Second, we train a long short-term memory (LSTM) network using Florida voter records for the general population, excluding drivers on the Lyft platform. The LSTM network is then fine-tuned on data from the EIR respondents. Third, we train a Bayes classifier using public census data on common last names associated with geographic census blocks and race, as well as information about phone manufacturer, carrier and locale settings from EIR respondents. Finally, we use an XGBoost classifier that combines the predictions from first three models to produce a final race prediction. Overall our model achieves 87.6% accuracy.

1.6 Speed Limit Data

We obtain speed limit and other road characteristics information from FDOT's Open Data Hub. Our models include controls for speed limit, road lane count, and log average annual daily traffic.

2 Identification Strategies

For both identification strategies, we make estimation more computationally tractable by using weighted least squares after randomly sampling 0.05 percent of observations that are not associated with a citation, resulting in an analysis sample of nearly 20 million observations. Throughout we conduct inference with robust standard errors that are clustered by driver.

2.1 Fixed Effect Model

The specification of our fixed effect model (i.e., "FE Model") goes as follows. We control for time with year, month, day of week, and hour of day fixed effects. We account for geography with fixed effects for whether a driver was operating in a given geohash of approximately 1,200 by 600 meters. These controls allow us to account for effects driven by over-policing in certain locations. We also account for road features with measures of lane count and log average annual daily traffic. To control for features of a car, we use manufacturer and state of vehicle registration fixed effects, as well as controls for vehicle age and vehicle age squared. To account for features of a driver other than race, we use data observed on driver age, age squared, and gender. We also control for the different phases of driving on Lyft's platform with fixed effects for whether a driver was waiting to be offered a ride, driving to pick up a rider, or providing a ride. Supplementary Appendix Table S3 presents summary statistics for the continuous variables used as controls in the FE model by race and ethnicity of the Lyft drivers in our analysis sample.

2.2 Double Machine Learning Model

Our double machine learning model (i.e., "DML Model") uses a machine learning algorithm to automatically learn the specification through which the controls enter the regression model of citations. Therefore, relative to the FE model, the DML model can be thought of as implementing the same basic approach as the FE model, but with a more flexible and data-driven specification. All variables that we use for the FE model also enter into our DML model with two exceptions. First, the DML model uses raw coordinates of longitude and latitude rather than fixed effects for each geohash. Second, rather than using hour of day, day of week, and month fixed effects, we supply the ML model with the exact second of week and day of year at which each observation occurred. This allows the DML model to automatically partition space

and time based on the data, rather than relying on geohashes and hours, which are inherently arbitrary partitions.

Our DML results make use of two "black box" machine learning models, $\hat{f} : \mathbb{R}^K \to \mathbb{R}$, trained to predict either whether a driver is a minority or whether the ping results in a traffic stop based on vectors of ping-level observables $X^i = (x_1^i, \ldots, x_K^i)$. Both of these models take the form of ensembles of decision trees. To better understand their behavior, we use the method of SHAP values (*38*). SHAP values provide an additive decomposition, based on the theory of Shapley values, of each model prediction $\hat{f}(X^i)$ into components $\phi_k(X^i; \hat{f})$ "attributable" to each feature in the model, plus an intercept $\phi_0(X^i; \hat{f})$: $\hat{f}(X^i) = \phi_0(X^i; \hat{f}) + \sum_{k=1}^K \phi_k(X^i; \hat{f})$. A large value of $|\phi_k(X^i; \hat{f})|$ indicates that the kth feature has a large effect on the model's prediction $\hat{f}(X^i)$, while the sign of $\phi_k(X^i; \hat{f})$ indicates whether the value of the kth feature has a positive or a negative impact on the model's prediction. To construct the values $\phi_k(X^i; \hat{f})$, we use the TreeSHAP algorithm of (*39*).

The TreeSHAP algorithm begins by recursively defining $\hat{T}_S(X^i)$, the prediction of a decision tree $\hat{T}$ on the observation X^i when only a subset $S \subseteq \{1, \ldots, K\}$ of features is observed, as follows:

- If the tree $\hat{T}$ consists of only a sinlge node with no children, then $\hat{T}_S(X^i)$ is the prediction at this node.

- Suppose the top node of the tree splits on $x_k \leq t$. If $k \in S$, then $\hat{T}_S(X^i) = \hat{T}_S^L(X^i)$ if $X_k^i \leq t$ and $\hat{T}_S(X^i) = \hat{T}_S^R(X^i)$ otherwise, where $\hat{T}^L$ and $\hat{T}^R$ are the left and right subtrees. If $k \notin S$, then $\hat{T}_S(X^i) = w^L \hat{T}_S^L(X^i) + w^R \hat{T}_S^L(X^i)$, where w^L and w^R are the proportions of the training samples falling in the left or right children, respectively.

For models $\hat{f} = \sum_j \hat{T}^j$ that are ensembles of trees, define $\hat{f}_S(X^i) = \sum_j \hat{T}_S^j(X^i)$. The SHAP value $\phi_k(X^i; \hat{f})$ is then defined as the Shapley value corresponding to the kth feature in the pre-

diction $\hat{f}(X^i)$, where the prediction resulting from a subset S of features is defined as $\hat{f}_S(X^i)$:

$$\phi_k(X^i;\hat{f}) = \sum_{S\subseteq\{1,\ldots,K\}:k\notin S} \frac{|S|!(K-|S|-1)!}{K!}[\hat{f}_{S\cup\{k\}}(X^i) - \hat{f}_S(X^i)]$$

and $\phi_0(X^i;\hat{f}) = \hat{f}_\emptyset(X^i)$. In words, $\phi_k(X^i;\hat{f})$ is the difference between the model's prediction on X^i when the kth feature is observed and when it is not, averaged over all possible choices of which other features are observed. (*38*) show that this definition is the only one satisfying the properties of *local accuracy* ($\hat{f}(X^i) = \phi_0(X^i;\hat{f}) + \sum_{k=1}^K \phi_k(X^i;\hat{f})$), *consistency* (if $\hat{f}'_{S\cup\{k\}}(X^i) - \hat{f}'_S(X^i) \geq \hat{f}_{S\cup\{k\}}(X^i) - \hat{f}_S(X^i)$ for all S, then $\phi_k(X^i;\hat{f}') \geq \phi_k(X^i;\hat{f})$), and a technical condition called *missingness*. (*39*) provide a polynomial-time algorithm (TreeSHAP) for the computation of these values.

Supplementary Appendix Figures S1 and S2 show the distributions of SHAP values for each covariate for the minority- and citation-prediction models, respectively. Location and time covariates are among the most important, as measured by the average magnitude of SHAP values, for both models. For the minority-prediction model, other influential covariates include driver and vehicle age and vehicle make. For the citation-prediction model, the two most important covariates are the ping speed and the road's speed limit. Driver age and road traffic are also important, with high values of either tending to reduce the predicted probability of citation.

3 Regression Results

3.1 Citations

Supplementary Appendix Table S4 presents the estimated effect of minority status on citations reported in Figure 2 of our manuscript in regression table form.

3.2 Fines

Supplementary Appendix Table S5 presents the estimated effect of minority status on fines reported in Figure 2 of our manuscript in regression table form.

3.3 Accidents

Supplementary Appendix Table S6 presents the estimated effect of minority status on accidents and the estimated effect of minority status on accidents at the levels of speeding reported in Figure 3 of our manuscript in regression table form.

3.4 Reoffenses

Supplementary Appendix Table S7 presents the estimated effect of minority status on reoffending speed limits in each week relative to the receipt of a speeding ticket, which is plotted in Figure 4 of our manuscript in regression table form.

4 Robustness of Results

4.1 Results by Speeding

Supplementary Appendix Figures S4 and S5 respectively plot the effect of minority status on citations and fines for speeding at discrete levels of speeding.

4.2 Results by Sets of Controls in FE Model

Supplementary Appendix Table S8 presents estimated effects of minority status on citations by sets of controls. Column 1 presents estimates without any controls. Column 2 adds the fixed effects for different levels of exceeding the speed limit. Column 3 adds fixed effects for geohashes in the state of Florida and fixed effects for aspects of the road. Column 4 adds fixed effects for calendar time and the task a Lyft driver is completing (namely, waiting to be paired

with a passenger, picking up a passenger, or providing a trip). Column 5 adds controls for features of the driver and their car. Across these specifications, we see that without controls minority drivers appear to be cited less frequently than white drivers. Adding controls for driving speeds magnifies the effect in column 1, however adding the location and road controls in column 3 flips the sign to indicate minorities are more likely to be cited for speeding. The time controls added in column 4 have a negligible effect on the estimates, whereas the driver and car controls in column 5 shrink the magnitude of the estimate slightly.

4.3 Results by Inference of Timing of Stop

A straightforward concern with our estimates is that our strategy for inferring when a ping led to a driver being cited for speeding is inaccurate. In the paper, we search for the largest reduction in driving speed within 3 minutes of the time cited in the Florida records and assign the ping that preceded this slowdown to the citation. In Supplementary Appendix Tables S9 and S10 we consider alternative thresholds of 1 and 5 minutes. Across each model, outcome, and threshold we see that our qualitative results are unaltered suggesting that our strategy for inferring the pings that led to a driver being cited is not a determinant of our results.

4.4 Results by Race or Ethnicity

Many are interested in the effect of racial profiling on specific racial and ethnic groups. While the total number of citations in our analysis sample is small and statistical precision for finer cuts than minority relative to white motorists is limited, we nonetheless present estimates for each racial and ethnic group in column 1 of Supplementary Appendix Tables S11 and S12. Therein, we see that the coefficient on Asian and Pacific Islander (API) is the largest in magnitude but estimated relatively imprecisely. The coefficient on Hispanic is the second largest coefficient and the mostly precisely estimated, while the coefficient on Black is the smallest in magnitude.

4.5 Results with Alternative Inference of Driver Race or Ethnicity

To establish the robustness of our results to alternative procedures of inferring race, we infer race using the frequency of race and ethnicity by surname from the 2020 Census. Column 2 in Supplementary Appendix Table S11 and S12 reports the effect of minority status on citations with race inferred from the Census. Nearly 8 million observations are dropped in this column because we drop drivers whose surname does not appear in the Census frequency tables. Column 3 reports the effect of minority status supposing that the drivers whose surnames do not appear in the Census frequency tables belong to a racial or ethnic minority. Across the two specifications we see that the drivers dropped in column 2 shrink the effect of minority status, while the effect of minority on fines remains statistically significant at the five percent level, the effect is marginally significant on citations with a p-value of 0.058. Moving to column 3, we see that the effect increases slightly when the drivers dropped in column 2 are assumed to belong to a racial or ethnic minority. Broadly speaking these results suggest that our results are robust to alternative approaches of inferring driver race or ethnicity.

4.6 Results with Post-Covid and Post-Floyd Omitted

The end of our sample was a relatively extraordinary time in the history of the United States, with lockdowns in response to the Covid-19 pandemic after March 2020 and large-scale unrest in response to the killing of George Floyd after May 2020. To establish that these extraordinary moments in history are not significantly altering our estimates, we presented estimated effects of minority status on citations omitting these two windows of time. Results are presented in Supplementary Appendix Table S13, which show that dropping these windows of time only alters the third decimal place of our estimated effects.

4.7 Results by Subsamples of Uncited Pings

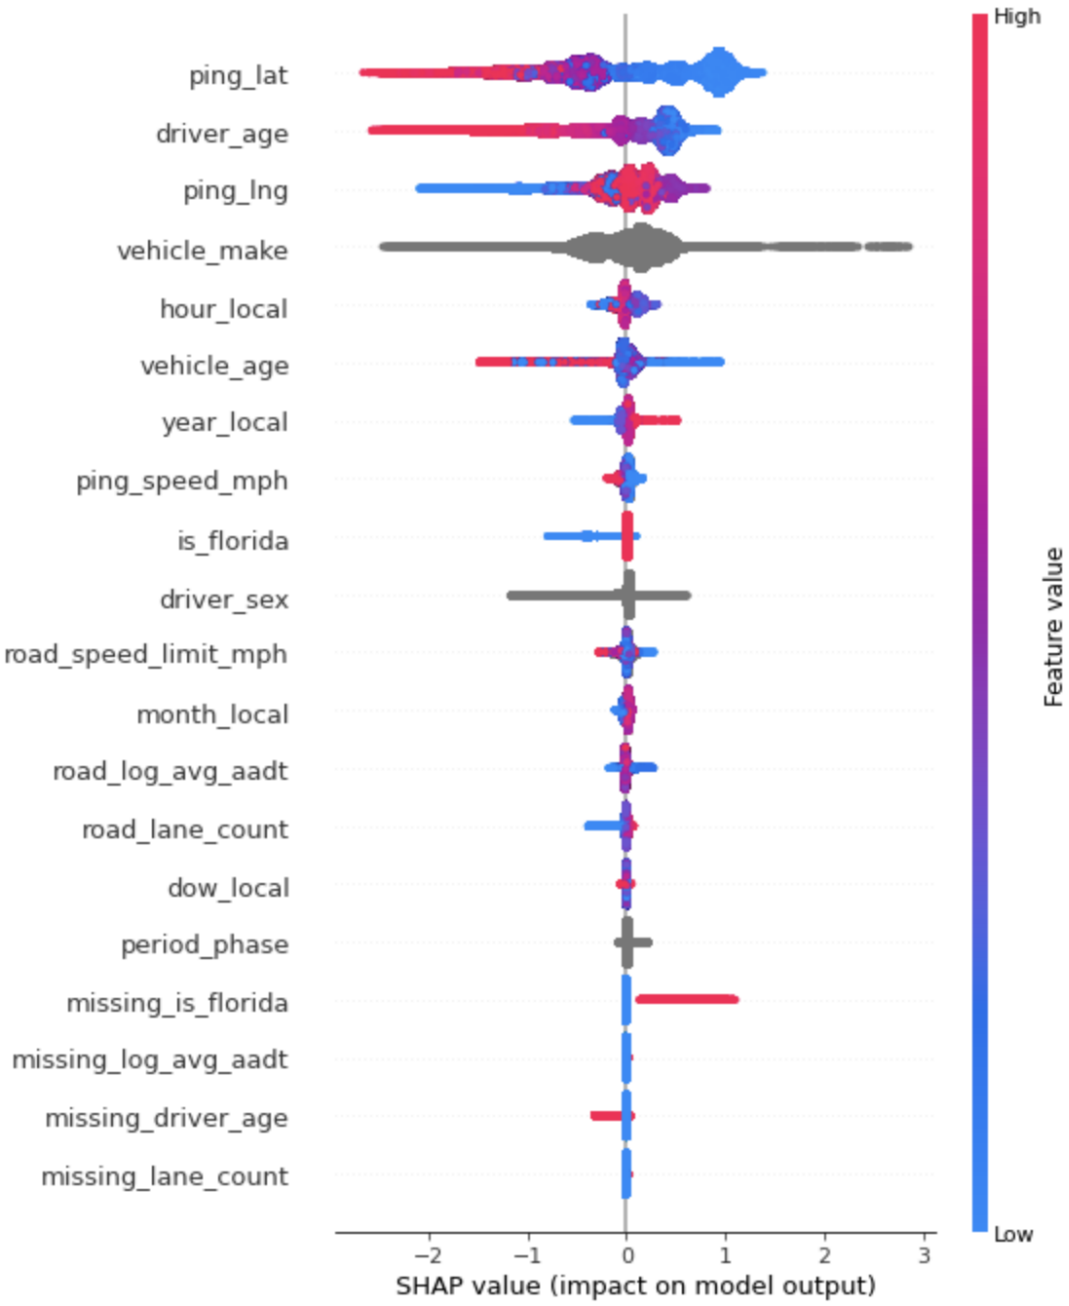
To evaluate whether a specific subsample of uncited pings are driving our results, we break the uncited pings into randomly selected quintiles and estimate our FE model. Supplementary Appendix Table S14 and S15 presents the results, with estimated effects all statistically significant at the one percent level.

References

38. S. M. Lundberg, S.-I. Lee, *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17 (Curran Associates Inc., Red Hook, NY, USA, 2017), pp. 4768–4777.

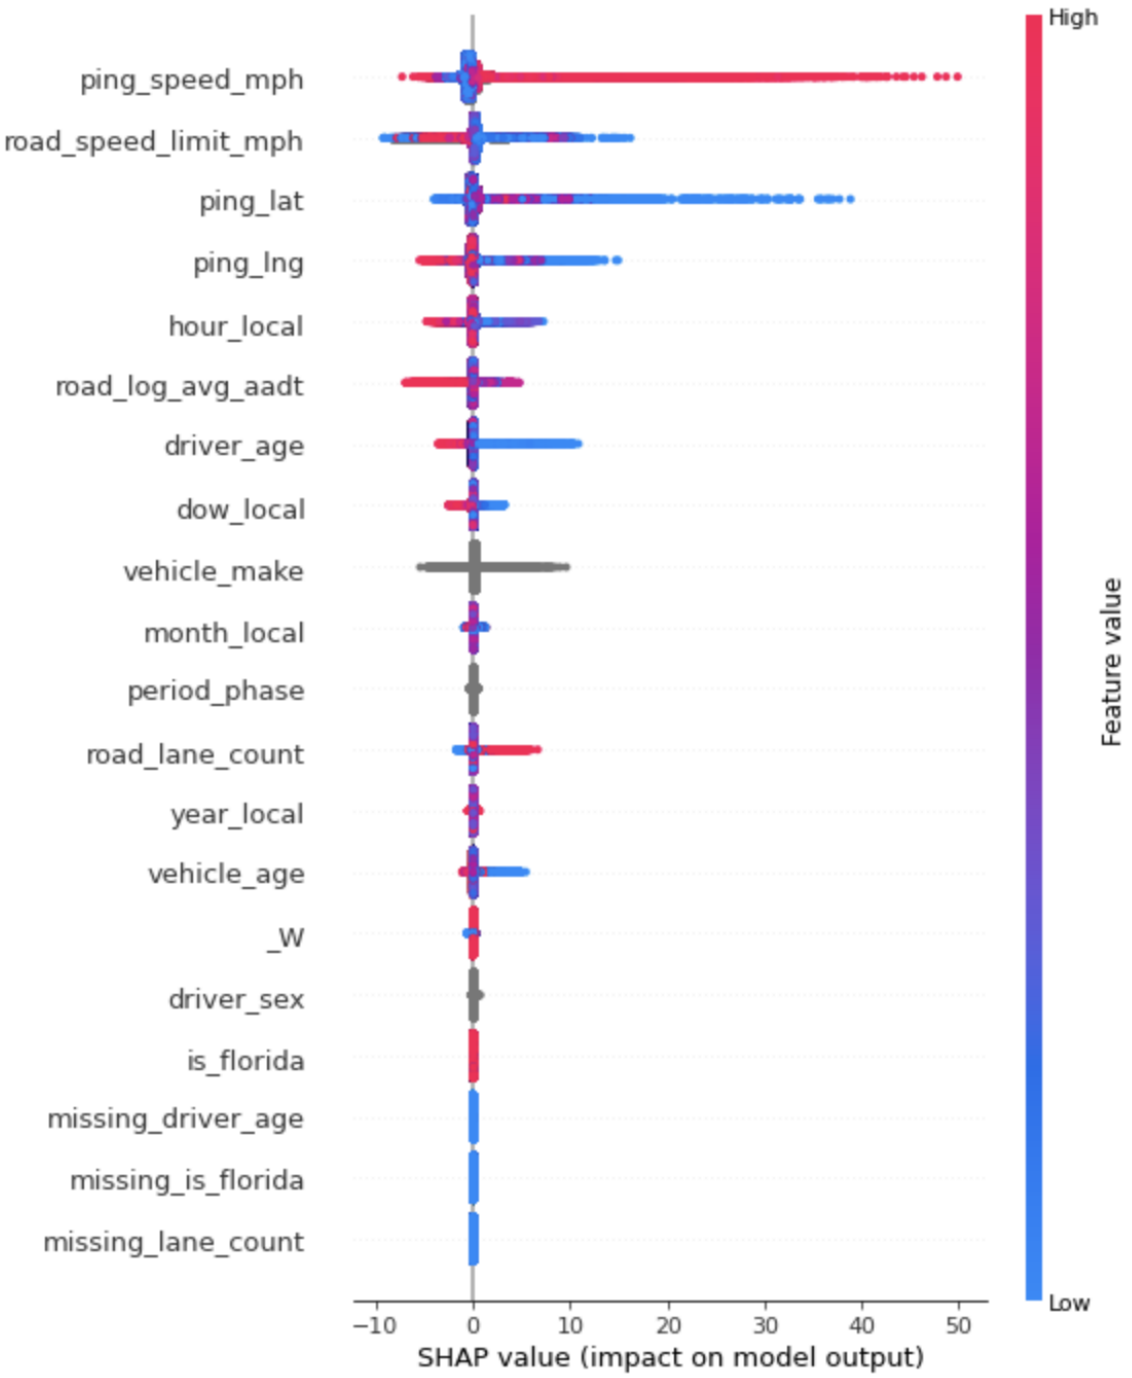
39. S. M. Lundberg, G. G. Erion, S.-I. Lee, *arXiv preprint arXiv:1802.03888* (2018).

Figure S1
SHAP values for predicting driver race



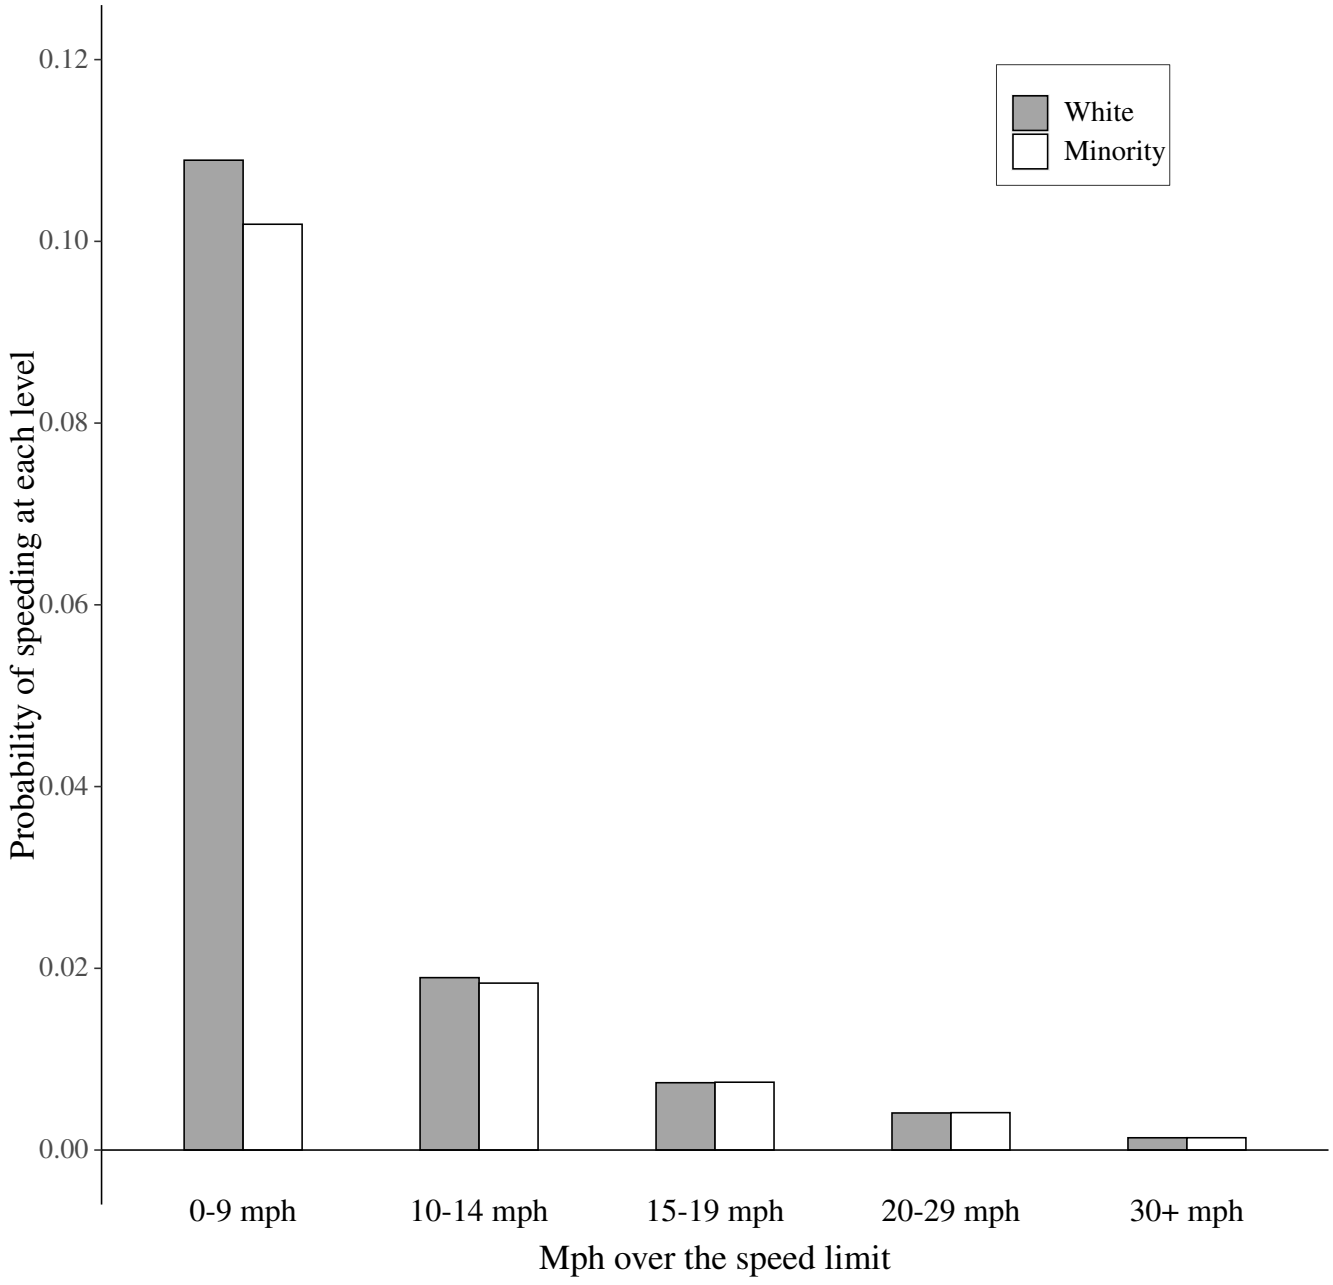
Note: This figure plots the distribution of SHAP values for each covariate for the minority-prediction model.

Figure S2
SHAP values for predicting citations



Note: This figure plots the distribution of SHAP values for each covariate for the citations-prediction model.

Figure S3
Speeding Behavior of Lyft Drivers in Florida by Race Group



Note: This figure plots the proportion of pings transmitted to Lyft in different speeding buckets. Speeding buckets are constructed by comparing the driving speed received by Lyft to the speed limit reported by the FDOT.

Figure S4
Effect of driver race on citations by speeding bucket

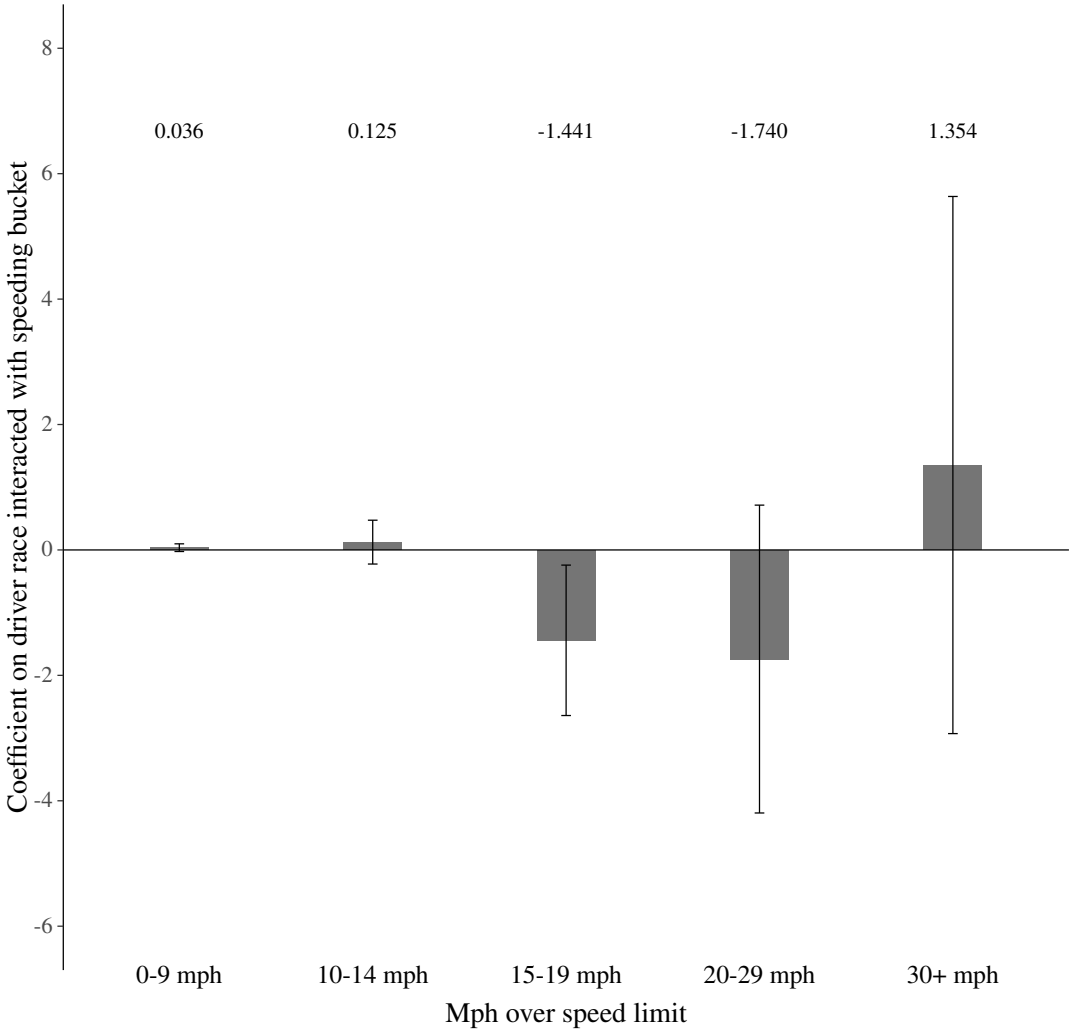


Figure S5
Effect of driver race on fines by speeding bucket

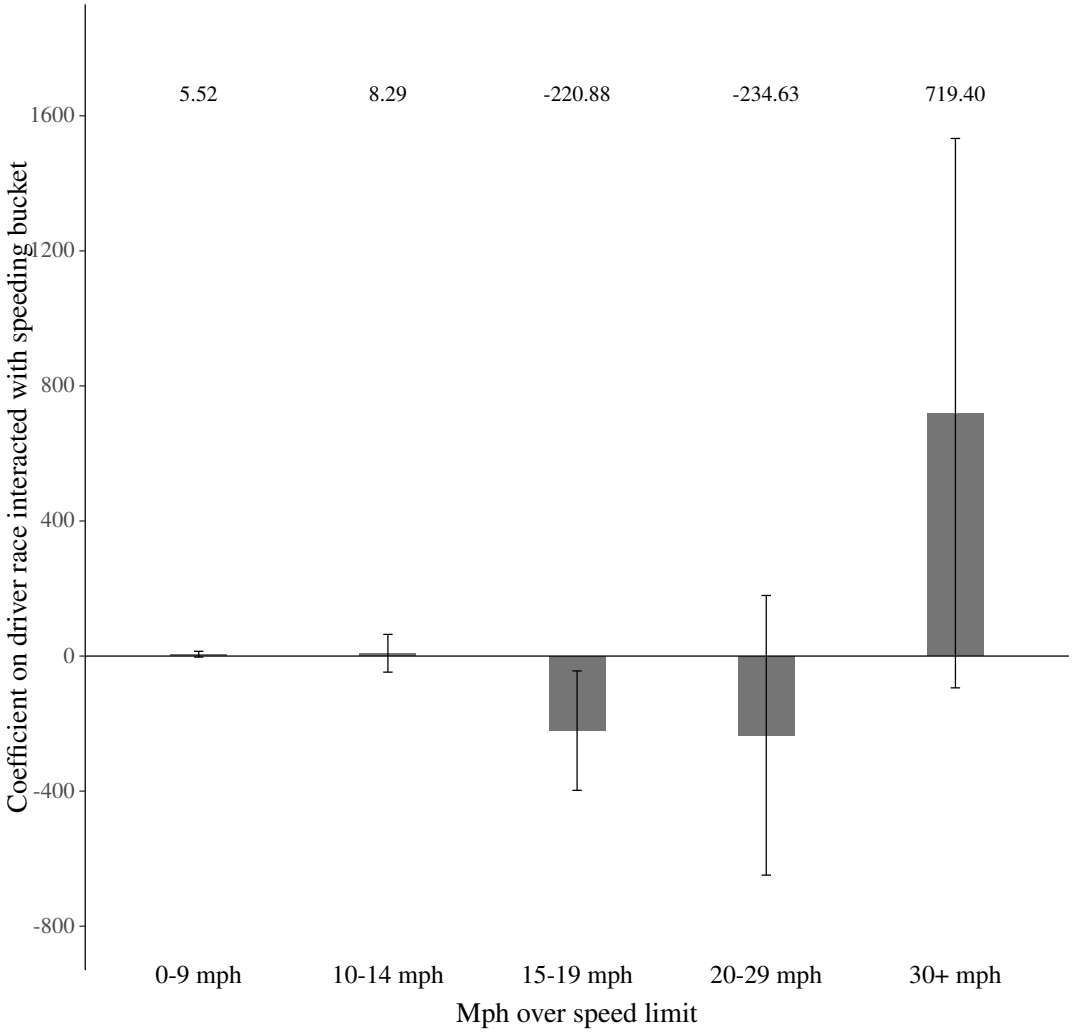


Table S1
Summary statistics on speeding behavior

	White	Black	Hispanic	API	Minority	Total
Number of drivers	61,510	56,134	101,240	3,954	161,328	222,838
	(27.6%)	(25.2%)	(45.4%)	(1.77%)	(72.4%)	
Average speed relative to speed limit	−22.8	−22.5	−23.9	−24.1	−23.4	−23.2
	(14.4)	(28.6)	(12.6)	(9.67)	(19.7)	(18.4)
Average portion of time not speeding	0.836	0.842	0.868	0.856	0.859	0.853
	(0.163)	(0.154)	(0.133)	(0.149)	(0.141)	(0.148)
Average portion of time 0-9mph over speed limit	0.130	0.119	0.100	0.111	0.107	0.113
	(0.145	(0.134)	(116)	(0.127)	(0.123)	(0.130)
Average portion of time 10-14mph over speed limit	0.021	0.023	0.018	0.020	0.020	0.020
	(0.059)	(0.060)	(0.049)	(0.063)	(0.053)	(0.055)
Average portion of time 15-19mph over speed limit	0.007	0.009	0.008	0.007	0.008	0.008
	(0.034)	(0.039)	(0.031)	(0.029)	(0.034)	(0.034)
Average portion of time 20-29mph over speed limit	0.004	0.005	0.004	0.004	0.005	0.004
	(0.028)	(0.027)	(0.025)	(0.025)	(0.026)	(0.026)
Average portion of time 30+mph over speed limit	0.002	0.002	0.001	0.002	0.002	0.002
	(0.024)	(0.020)	(0.016)	(0.024)	(0.018)	(0.020)
Average citations per 10,000 speeding hours	0.141	0.158	0.125	0.177	0.136	0.136
Average fines per 10,000 speeding hours	19.1	20.8	18.5	25.2	19.3	19.2

Note: Standard deviations in parentheses

Table S2
Sample characteristics

	2017 FL population aged 18+	Cited FL drivers	Lyft drivers	Cited Lyft drivers
Portion female	0.517	0.414	0.289	0.198
	(0.500)	(0.493)	(0.453)	(0.398)
Average age	47.7	37.5	40.4	40.4
		(15.1)	(12.6)	(12.3)
Portion White	0.549	0.665	0.276	0.253
	(0.498)	(0.472)	(0.447)	(0.435)
Portion Black	0.154	0.183	0.252	0.242
	(0.361)	(0.387)	(0.434)	(0.428)
Portion Hispanic	0.247	0.060	0.454	0.482
	(0.431)	(0.237)	(0.498)	(0.500)
Portion Other	0.050	0.093	0.018	0.023
	(0.218)	(0.290)	(0.133)	(0.150)
Portion minority	0.451	0.335	0.724	0.748
	(0.498)	(0.472)	(0.447)	(0.434)
Number of drivers	12,336,038	7,172,871	222,833	1,408

Note: Portion of each race for cited FL drivers is based on the police's definition of driver race.
Standard deviation in parentheses.

Table S3
Summary Statistics on control variables

	White	Black	Hispanic	API	Minority	Total
Number of drivers	61,510 (27.6%)	56,134 (25.2%)	101,240 (45.4%)	3,954 (1.77%)	161,328 (72.4%)	222,838
Average portion of P1 time	0.519 (0.270)	0.459 (0.261)	0.525 (0.241)	0.511 (0.272)	0.502 (0.251)	0.507 (0.256)
Average portion of P2 time	0.144 (0.156)	0.158 (0.159)	0.118 (0.132)	0.146 (0.163)	0.133 (0.144)	0.136 (0.148)
Average portion of P3 time	0.336 (0.228)	0.382 (0.227)	0.355 (0.210)	0.342 (0.230)	0.364 (0.217)	0.356 (0.221)
Portion of female drivers	0.283	0.387	0.244	0.122	0.291	0.289
Average driver age	44.0 (14.2)	37.5 (11.5)	39.8 (11.5)	40.8 (12.2)	39.0 (11.6)	40.4 (12.6)
Portion of FL registration	0.906	0.929	0.967	0.921	0.953	0.940
Average road lane count	2.31 (0.385)	2.37 (0.380)	2.41 (0.341)	2.37 (0.388)	2.40 (0.357)	2.37 (0.367)
Average log avg AADT	10.0 (0.504)	10.1 (0.470)	10.1 (0.416)	10.1 (0.480)	10.1 (0.438)	10.1 (0.458)
Average vehicle age	5.29 (3.56)	5.46 (3.46)	4.65 (3.29)	5.20 (3.52)	4.94 (3.38)	5.04 (3.43)

Note: Missing observations are omitted for the respective outcomes
Standard deviations in parentheses

Table S4 Citation Results

	Citations per 10,000 Hours of Driving Time	
	(1)	(2)
Minority	0.036	0.027
	$(0.011)^{***}$	$(0.010)^{***}$
Model	FE	DML
Pings	19,356,683	19,356,683

*** p-value < 0.01, ** p-value < 0.05, * p-value < 0.10.

Table S5 Fine Results

	Fines in USD per 10,000 Hours of Driving Time	
	(1)	(2)
Minority	5.20	3.54
	$(1.73)^{***}$	$(1.56)^{**}$
Model	FE	DML
Pings	19,356,683	19,356,683

*** p-value < 0.01, ** p-value < 0.05, * p-value < 0.10.

Table S6 Accident Results

	Accidents per 10,000 Hours of Driving Time	
	(1)	(2)
Minority	0.09	0.07
	(0.08)	(0.07)
Minority X 0-9mph		0.22
		(0.14)
Minority X 10-14mph		−0.03
		(0.32)
Minority X 15-19mph		−0.49
		(0.55)
Minority X 20-29mph		0.18
		(0.70)
Minority X 30mph+		−2.70
		(2.30)
Model	FE	FE
Pings	19,356,683	19,356,683

*** p-value < 0.01, ** p-value < 0.05, * p-value < 0.10.

Table S7 Re-Offense Result

	Proportion of Time Driving 10mph Over Speed Limit
	(1)
Minority X -7 Weeks	-0.004
	(0.003)
Minority X -6 Weeks	0.001
	(0.003)
Minority X -5 Weeks	0.000
	(0.003)
Minority X -4 Weeks	0.002
	(0.003)
Minority X -3 Weeks	0.001
	(0.002)
Minority X -2 Weeks	0.002
	(0.003)
Minority X -1 Week	0.001
	(0.002)
Minority X 1 Week	-0.001
	(0.002)
Minority X 2 Weeks	0.001
	(0.002)
Minority X 3 Weeks	0.000
	(0.003)
Minority X 4 Weeks	0.002
	(0.002)
Minority X 5 Weeks	0.002
	(0.003)
Minority X 6 Weeks	0.004
	(0.003)
Minority X 7 Weeks	0.003
	(0.003)
Minority X 8 Weeks	0.000
	(0.003)
Observations	22,768

Note: This table presents the coefficient of minority status interacted with weeks relative to a driver receiving a speeding citation. The unit of observation is a week relative to a driver receiving a speed citation. Coefficients are obtained by regressing the portion of time in a week spent driving more than 10 miles per hour over the speed limit on the interaction between minority status and weeks since a citation for speeding, driver fixed effects, and weeks since a citation fixed effects. *** p-value < 0.01, ** p-value < 0.05, * p-value < 0.10.

Table S8 Results from FE Model with Controls Added Sequentially

	Citations per 10,000 Hours of Driving Time				
	(1)	(2)	(3)	(4)	(5)
Minority	-0.005	-0.015	0.056	0.054	0.036
	(0.009)	$(0.009)^{*}$	$(0.011)^{***}$	$(0.011)^{***}$	$(0.011)^{***}$
Controls	None	Speeding	Speeding, Location, Road	Speeding, Location, Road, Time	Speeding, Location, Road, Time, Driver, Car
Pings			19,356,683		
RMSE	31,382.9	31,382.9	31,361.5	31,361.5	31,361.5

*** p-value < 0.01, ** p-value < 0.05, * p-value < 0.10.

Table S9 Robustness check with altered definition of cited ping (FE Model)

Threshold for cited ping		1 minute	3 minutes	5 minutes
Citations	Average citations per 10,000 hours for white drivers	0.108	0.110	0.110
	Average citations per 10,000 hours for minority drivers	0.144	0.146	0.146
	Difference between race groups	0.037** (33.9%)	0.036** (32.8%)	0.037** (33.3%)
Fines	Average fine per 10,000 hours for white drivers	14.9447	15.286	15.217
	Average fine per 10,000 hours for minority drivers	20.382	20.486	20.598
	Difference between race groups	5.437** (36.4%)	5.201** (34.0%)	5.381** (35.4%)

*** p-value < 0.01, ** p-value < 0.05, * p-value < 0.10.

Table S10 Robustness check with altered definition of cited ping (DML Model)

Threshold for cited ping		1 minute	3 minutes	5 minutes
Citations	Average citations per 10,000 hours for white drivers	0.112	0.110	0.112
	Avg citations per 10,000 hours for minority drivers	0.138	0.138	0.138
	Difference between race groups	0.026** (23.5%)	0.027** (24.4%)	0.026** (23.2%)
Fines	Average fine per 10,000 hours for white drivers	15.744	15.763	15.662
	Average fine per 10,000 hours for minority drivers	19.223	19.344	19.346
	Difference between race groups	3.479** (22.1%)	3.581** (22.7%)	3.684** (23.5%)

*** p-value < 0.01, ** p-value < 0.05, * p-value < 0.10.

Table S11 FE Model Results by Race or Ethnicity and Alternative Inference of Minority

	Citations per 10,000 Hours of Driving Time		
	(1)	(2)	(3)
API	0.056		
	$(0.034)^{*}$		
Black	0.023		
	$(0.013)^{*}$		
Hispanic	0.042		
	$(0.012)^{***}$		
Alt. Minority		0.022	0.033
		$(0.011)^{*}$	$(0.010)^{**}$
Pings	19,356,683	11,575,959	19,356,683

*** p-value < 0.01, ** p-value < 0.05, * p-value < 0.10.

Table S12 FE Model Results by Race or Ethnicity and Alternative Inference of Minority

	Fines per 10,000 Hours of Driving Time		
	(1)	(2)	(3)
API	9.847		
	$(5.607)^{*}$		
Black	2.765		
	(2.023)		
Hispanic	6.375		
	$(1.907)^{***}$		
Alt. Minority		4.749	6.236
		$(1.831)^{***}$	$(1.550)^{***}$
Pings	19,356,683	11,575,959	19,356,683

*** p-value < 0.01, ** p-value < 0.05, * p-value < 0.10.

Table S13 FE Model Results Omitting Post-March 2020 and Post-May 2020

	Citations per 10,000 Hours of Driving Time	
	(1)	(2)
Minority	0.030	0.032
	$(0.012)^{**}$	$(0.011)^{**}$
Dates Omitted	Post-March 2020	Post-May 2020
Pings	17,761,276	18,407,496

*** p-value < 0.01, ** p-value < 0.05, * p-value < 0.10.

Table S14 Results from FE Model by Subsample of Uncited Pings

	Citations per 10,000 Hours of Driving Time				
	(1)	(2)	(3)	(4)	(5)
Minority	0.037	0.041	0.038	0.034	0.033
	$(0.011)^{***}$	$(0.011)^{***}$	$(0.011)^{***}$	$(0.011)^{***}$	$(0.011)^{***}$
Pings	3,872,514	3,872,514	3,872,513	3,872,513	3,872,513

*** p-value < 0.01, ** p-value < 0.05, * p-value < 0.10.

Table S15 Results from FE Model by Subsample of Uncited Pings

	Fines per 10,000 Hours of Driving Time				
	(1)	(2)	(3)	(4)	(5)
Minority	5.471	6.215	5.502	5.052	4.930
	$(1.821)^{***}$	$(1.798)^{***}$	$(1.806)^{***}$	$(1.778)^{***}$	$(1.805)^{***}$
Pings	3,872,514	3,872,514	3,872,513	3,872,513	3,872,513

*** p-value < 0.01, ** p-value < 0.05, * p-value < 0.10.