Walking while Black : Racial Gaps in Hit-and-Run Cases

Madina Kurmangaliyeva
: ECARES, Université libre de Bruxelles & TILEC, Tilburg University

Matteo Sostero
Joint Research Centre, European Commission

April 2022

ECARES working paper 2022-08

ECARES
ULB - CP 114/04
50, F.D. Roosevelt Ave., B-1050 Brussels BELGIUM
www.ecares.org

Walking While Black: Racial Gaps in Hit-and-Run Cases[*]

Madina Kurmangaliyeva[**]   Matteo Sostero[†]

April 5, 2022

Abstract

We provide a causal test for racial gaps in victimization and clearance rates, using unintentional vehicle-pedestrian crashes. The victim's race should not depend on the driver's characteristics, conditional on location and time. We find that American drivers flee 13% more often if they hit Black pedestrians, and their clearance rates are 11% lower. This provides rare evidence of racial discrimination by the public in a high-stakes environment. These gaps correlate, suggesting statistical discrimination as a mechanism and underlining the importance of closing the racial gap in clearance rates, especially in poorer non-Black neighborhoods. Taste-based discrimination is arguably also at play.

JEL codes: *J15, K40, K42*

Keywords: *racial gap, victimization, hit-and-run crimes, statistical discrimination, out-group bias, law enforcement*

[*]We would like to thank Andrea Ichino, Andrea Mattozzi, Juan Dolado, Roberto Galbiati, Ben Vollaard, Joseph Doyle, Ada González-Torres, Gani Aldashev, Cedric Argenton, David Schindler, Amanda Agan, Steve Mello, Sarit Weisburd, Jennifer Doleac, participants at the APPAM 2020, the Texas Economics of Crime Workshop 2020, Spring Meeting of Young Economists 2019, CEPR Women in Economics Workshop in Micro- & Applied Microeconomics, seminar participants at the European University Institute, Structural Econometrics Group (Tilburg), and Tilburg Economics Workshop. Special thanks to Jeremy West for pointing us to the CRIS database.

[**]Corresponding author: ECARES, Université libre de Bruxelles (ULB); TILEC (Tilburg University); madina.kurmangaliyeva@ulb.be

[†]European Commission - Joint Research Centre; matteo.sostero@ec.europa.eu. The content of this article does not reflect the official opinion of the European Union. Responsibility for the information and views expressed in therein lies entirely with the authors.

1 Introduction

Black Americans are more likely than white Americans to become victims of crime, and their cases are less likely to be cleared by the police (Harvey and Mattia 2019; Fagan and Geller 2018; Langley and Sugarmann 2017; Lowery et al. 2018; Lee 2005). While empirical studies have mostly focused on the differential treatment of Black defendants in the U.S. law enforcement and criminal justice systems (e.g., West 2015; Alesina and La Ferrara 2014; Persico and Todd 2006; Knowles et al. 2001), much less is known about treatment of victims based on their race, especially at the earlier stages of crime victimization.

Finding that offenders are more likely to commit crime when the potential victim is Black, rather than white, would have several important implications. First, measuring racial gaps in court outcomes would likely understate the full extent of racial disparity in the American law enforcement and criminal justice system, as this would not account for the disparate treatment in decisions happening upstream. Second, it raises the question of whether the racial disparity observed downstream of the criminal justice pipeline feeds back into higher victimization rates for the Black population. If crimes against Black victims are less likely to be cleared/prosecuted, are offenders more likely to commit crimes against those victims in the first place? If so, then the disparities in criminal justice would fail Black victims *twice*: directly, through unfair decisions of institutional actors to extant victims, and, indirectly, through under-deterrence of future crimes against Black victims.

In this paper, we test whether the racial gaps in victimization (and in clearance rates) of hit-and-run crimes can be causally attributed to the race of the victim. Specifically, we compare outcomes for traffic collisions involving severe injuries to Black and non-Hispanic white pedestrians in the same location.

In general, it is difficult to estimate how offenders' behavior would change if the victim's race had been different, since many crimes occur between people who know each other or happen to live in the same area: e.g., domestic violence, or the majority of murders. In other words, this endogenous relationship between offenders and victims makes it difficult to systematically compare cases across races. We argue that traffic collisions involving severe injuries to pedestrians are an appropriate setting to test for the equal treatment of victims. Firstly, traffic accidents with severe injuries are systematically recorded, which removes reporting bias. Also, drivers usually do not choose whom to hit in an accident (and the accident itself is not necessarily a crime), but the drivers do choose whether to stay or run, where failure to help an injured person usually constitutes a crime. Hence, the unintentional nature of traffic accidents alleviates concerns about most confounders in the relationship between offenders and victims, by bringing together complete strangers.[1] Moreover, while reckless and careless drivers are likely over-represented in the population of drivers at risk of an accident, no driver is completely safe from a traffic collision. Hence, this setting is as close as we can get to study racial discrimination by the general public in high-stakes environments.

Our main source of data is the Fatality Analysis Reporting System (FARS) for 2010–2016, which covers all car crashes in the U.S. that resulted in human deaths. We also use the TxDOT's Crash Records Information System (CRIS), which records all crashes, fatal or not, for the State of Texas.

Using FARS data, we restrict our attention to crashes involving a single vehicle and a single pedestrian, which represent roughly 85% of all accidents with pedestrian fatalities.[2] On average, there are 2,640 fatal cases per year in the U.S. involving either a white or a Black pedestrian. Raw numbers suggest that hit-and-runs are more prevalent for Black pedestrians (25%) than for white pedestrians (16%). At the same time, the rate at which the driver is eventually identified is higher for white pedestrians (48%) than for Black pedestrians (35%). However, we cannot conclude that drivers behave differently depending on the race of the pedestrian based on these numbers alone, since the crashes involving Black pedestrians tend to happen in different geographical areas than those involving white pedestrians. Hence, these cases are not directly comparable, as the population of drivers at risk is different for Black and white pedestrians.

Our strategy is to compare outcomes for collisions that happened in the same location. In practice, we define a location by the boundaries of a census tract or a cluster of several census tracts with homogeneous racial composition of residents. Crucially, the data contains the coordinates of the crash, other

1. This setting is also used by Glaeser and Sacerdote (2003) and Kurmangaliyeva (2018). Grogger and Ridgeway (2006), Horrace and Rohlin (2016), and West (2015) use other sources of randomization related to traffic stops and traffic offenses.

2. We do not focus on vehicle-vehicle collisions as in this setting it may not be clear who among the drivers is injured and who, if any, is in a position to flee. At the same time the injuries are endogenous with respect to the relative weight and safety of the vehicles.

crash characteristics, and demographic profiles of the victims and drivers (if they stayed or were identified). Race and ethnicity records are available for victims, but not for drivers.

Given the unintentional nature of traffic accidents, we assume that the identity of the pedestrian in a traffic collision is a random draw from the population of pedestrians at risk in a given location. This identifying assumption is similar to Levitt and Porter (2001) who assume "equal mixing" of drivers on roads. The probability of hitting a pedestrian of a certain race and the characteristics of drivers may thus change from neighborhood to neighborhood, but conditional on a given location, time, and circumstances before the crash, the characteristics of the driver are assumed to be independent of the race of the pedestrian. Importantly, if drivers do not discriminate based on the victim's race, then the local incidence of hit-and-runs should be the same for Black and white pedestrians.[3] Moreover, the victim's race should not correlate with the characteristics of the drivers who stayed. If, however, drivers are more likely to flee when the pedestrian is Black, then the frequency of hit-and-runs will be higher for Black pedestrians, and we will observe a *racial gap* in outcomes.

Our results show that drivers are 13% (21% when using stricter location controls) more likely to flee after hitting a Black pedestrian rather than a white pedestrian, while controlling for location fixed effects, hour of the day, and the characteristics of pedestrians (e.g., intoxication), crash (e.g., crosswalk, light conditions), and census block groups (e.g., income and ethnicity of local residents). A racial gap in hit-and-run rates seem to exist at any light condition, but it is especially pronounced at night-time with artificial lighting. The racial gap is larger for crashes occurring on local roads in poorer, whiter residential areas (130% more likely), or where Black residents are in minority (40% more likely). Only for crashes occurring on local roads in Black neighborhoods, the point estimate of the gap in hit-and-runs appears to be negative – i.e., drivers seem to discriminate against white victims – but it is not statistically significant. Our main results are robust to different definitions of location fixed effects, or to different functional forms, such as the inclusion of interaction terms. We also confirm that the racial composition of neighborhoods is the best predic-

3. The term *discrimination* is used in its literal meaning: "to make a distinction". It is an empirical statement, and does not necessarily imply a value judgment on the preferences of the drivers.

tor of heterogeneity in racial gaps, using a generalized random forest (Athey et al. 2019).

Given that FARS data only includes data for fatal crashes, we extend the analysis on the CRIS dataset that also records non-fatal crashes (but only for Texas). We show that our main estimates based on fatalities, if anything, understate the magnitude of the racial gap. We also find that race, rather than income, is the source of discrimination: in CRIS the racial gap in hit-and-runs remains almost unchanged when controlling for the income of pedestrians, as proxied by the income in their ZIP code. Likewise, residency does not seem to be driving the results, as the racial gap in hit-and-run rates in non-Black neighborhoods remain unchanged when we control for the residence status of the pedestrian.

Importantly, using FARS data, we also find tentative evidence that hit-and-run cases are 12% less likely to be cleared (i.e., the driver is identified) when the victim is Black rather than white, conditional on location, time, and crash characteristics. Again, this racial gap is wider in poorer areas where Black residents are in minority. This 12% gap in clearance rates cannot be attributed solely to the difference between drivers who fled for white victims and those who fled for Black victims: given the magnitude of the racial gap in hit-and-run rates, this would require that *all* marginal hit-and-run drivers – those who flee for Black pedestrians, but stay for white pedestrians – were impossible to find. Such a scenario seems implausible: the average clearance rate is at 44%. More likely, this racial gap stems from other reasons, such as differences in police efforts, or the resources and cooperation mobilized by the family of the victim and the local community.

Why would the decisions of drivers to flee change based on the race of the victim? Becker (1974)'s model of crime and punishment suggests that the drivers change their decisions because their expectations of legal sanctions change with the victim's race – i.e., statistical discrimination (Arrow et al. 1973) – or because they on average attach lower utility to helping injured Black pedestrians – i.e., taste-based discrimination (Becker 1971), or both. We explore both hypotheses in the rest of the paper.

It is conceivable that the unequal clearance rates could affect drivers' expectations of being caught, and could ultimately be the reason why Black pedestrians are more likely to become victims of a hit-and-run in the first place.

In other words, the systematic differences in the treatment of victims in the downstream of the criminal justice pipeline is affecting the upstream decisions of the offenders. Other papers studying hit-and-runs find that drivers do react rationally to changes in the expected punishment, or the probability of being caught (French and Gumus 2015; Castriota and Tonin 2019).

As expected, we find a negative relationship between predicted hit-and-run rates and the predicted case clearance rates (See Figure 3a on page 33), just as Becker (1974)'s model predicts. In line with the *statistical discrimination* hypothesis, we also observe that the settings where clearance rates are the most unequal for Black victims are also those where drivers discriminate the most against Black pedestrians (See Figure 3b). In other words, racial gaps in clearance rates may potentially explain racial gaps in hit-and-run rates.

However, even where we estimate no racial gap in clearance rates we still estimate a substantial racial gap in hit-and-run rates. When clearance rates are the same for Black and white victims, perfectly-informed drivers should not discriminate statistically based on clearance rates. Hence, either the drivers have biased beliefs about the racial gaps in clearance rates, or their decisions are also influenced by *taste-based discrimination* or expected racial disparities in later stages of law enforcement and criminal justice (i.e., fines, prison terms, etc.). At the same time, Glaeser and Sacerdote (2003) show that sentences are in fact milder for vehicular homicide cases when the victim is Black, which should make perfectly-informed drivers statistically discriminate against white pedestrians, not against Black pedestrians.

Indeed, we also find tentative evidence in support of *taste-based discrimination.* Conditional on location and time, the racial profile of drivers should be balanced across white and Black pedestrians. If Black drivers are more likely to stay for Black pedestrians, while white drivers are more likely to stay for white pedestrians, then the race of the hit-and-stay driver will be correlated with the race of the pedestrian. We find that this is the case, although only weakly significant, when using the stricter location controls). Hence, the driver's decision to flee may, in principle, also depend on the alignment of their own race with the race of the victim.

Our main results rely on the assumption that Black and white pedestrians, conditional on location and time, are hit by the same type of drivers (in expectation). This "equal mixing" assumption should also apply to other characteristics of pedestrians, not just race. We can thus apply the same test to discover whether drivers discriminate among pedestrians based on other observable characteristics. We find that drivers are more likely to stay for children and the elderly, compared to middle age groups. On the other hand, our test for equal treatment of victims show no difference for gender. While raw numbers suggest lower incidence of hit-and-runs for women (16.9%) than for men (18.4%), once we control for all the other characteristics of the crash, the share of hit-and-runs becomes balanced for female and male pedestrians, and so are their clearance rates. In other words, there appears to be no gender gap in hit-and-runs, while there is one for race and age.

Another important assumption of our empirical exercise is that Black and white pedestrians are hit in similar circumstances. In raw data, we see that the circumstances involving Black and white pedestrians differ substantially (light conditions, road type and its width, etc.). However, most of the observed differences disappear once we partial out location and time fixed effects, especially so when we use the narrowest definition of location, i.e., census tract fixed effects. Overall, all the additional controls added to location and time fixed effects, using the strictest definition of location, reduce the estimate of the racial gap, but only slightly: from 23.6% to 20.6%. The observable differences in crash circumstances (once we control for the census block group and pedestrian characteristics) have no additional effect on the estimated racial gap in hit-and-runs.

To our knowledge, this is the first paper that provides a test for racial discrimination at the earliest stage of the criminal justice pipeline, namely, the decision to victimize. Glaeser and Sacerdote (2003) show that there is a racial gap in sentencing of unintentional vehicular homicide cases: drivers who killed Black pedestrians receive much shorter sentences. However, this racial gap is only *conditional* on the case being cleared and prosecuted. Our paper shows that racial discrimination starts with the actions of the general public at the decision to victimize.

This finding has the following key policy implications. Law-enforcement authorities should monitor racial gaps in clearance rates, and attempt to increase clearance rates for Black victims even for non-intentional crimes, like hit-and-runs. Doing so not only has its own merit in terms of justice, but could also deter offenders in the first place. Several contributions focused on the interrelation between law enforcement and victimization. Harvey and Mattia (2019) find that more Black police officers in local police departments can close the racial gap in victimization rates. Comino et al. (2020) provide some evidence that giving legal protection to illegal immigrants in the U.S. can deter their victimization. Our study reveals that the most dramatic differences in clearance and hit-and-run rates are in poorer non-Black neighborhoods. Hence, the policy of monitoring these racial gaps can be targeted to those neighborhoods.

The existing literature on racial discrimination usually examines actions by the general public where the stakes are monetary (see review by Guryan and Charles 2013)., e.g., in employment decisions (e.g., Bertrand and Mullainathan 2004), charitable giving (e.g., Fong and Luttmer 2009), peer-to-peer lending (e.g., Duarte et al. 2015), housing market (e.g., Cutler et al. 1999). Or it looks at high-stakes decisions in institutional settings like capital sentencing (e.g., Alesina and La Ferrara 2014) or severity of incarceration policy (Feigenberg and Miller 2021). We find that racial discrimination also happens in high-stakes decisions made by the general public.

We also provide evidence that racial gaps persist at the investigation stage, contributing to the existing literature on the racial gaps in clearance rates: Fagan and Geller (2018) document racial gaps in clearance rates for capital homicide cases in the U.S. We show that racial gaps in clearance rates also exist for hit-and-run crimes. More broadly, this paper contributes to the literature on causal tests for racial (or ethnic) discrimination of victims in law enforcement and criminal justice (e.g., Glaeser and Sacerdote 2003; Shayo and Zussman 2011; Alesina and La Ferrara 2014).

Our paper is also tangentially related to the literature on racial bias in police decisions, e.g., in stop-and-frisk policy (e.g., Knowles et al. 2001; Anwar and Fang 2006; Grogger and Ridgeway 2006), in traffic citations (West 2015). This literature usually requires the identity of the police officer to be independent of the race of the driver. Instead, our paper tests for racial discrimination of pedestrians by drivers at risk of accidents. So our identification strategy relies on the "equal mixing" of pedestrians and drivers at risk in traffic collisions.

The next section provides detailed information on the sources of data used in the analysis.

2 Data

This paper uses data from the U.S. Fatality Analysis Reporting System (FARS), which covers all fatal crashes in each American State. Additionally, we use the data from the Crash Records Information System (CRIS), which is limited to the State of Texas, but includes all cases resulting in incapacitating injuries to pedestrians, whether fatal or not.

FARS contains information about all traffic accidents in the U.S. that resulted in at least one death within 30 days after the accident. The data is compiled by the National Highway Traffic Safety Administration, sourcing information from police reports, death certificates, state vehicle registration files, medical reports, and other documents from each state. The data comes as a set of yearly databases at the levels of crash, person, and vehicle, which are available under open access on the FARS website.[4]

The variable of interest is the hit-and-run variable "Hit_Run". It comes from a vehicle form and refers to "cases where a vehicle is a contact vehicle in the crash and does not stop to render aid (this can include drivers who flee the scene on foot)". If the driver of a vehicle is flagged as a hit-and-run driver, but there is full information about the driver (e.g., sex, age, ZIP Code, car information, etc.), we assume that the driver has been *identified* by the police.[5] FARS records the race of the pedestrian based on the information from the death certificate.

We restrict our sample to crashes involving a single driver who fatally hits a single pedestrian. In the period 2010 to 2016, the population of interest includes 18, 501 crashes for which coordinates and light conditions are recorded: 13, 710 non-Hispanic white pedestrians, and 4, 791 Black pedestrians.[6] Table 1 (Panel A) provides mean statistics for white and Black pedestrians across different

4. See *FARS: National Highway Traffic Safety Administration. Fatality Analysis Reporting System*: https://www.nhtsa.gov/research-data/fatality-analysis-reporting-system-fars

5. Note that we cannot distinguish between the runaway drivers that have been apprehended by the police and the runaway drivers who voluntarily turned themselves in later.

6. We exclude 290 crashes with missing longitude and latitude coordinates and 67 crashes with missing or unknown light conditions on roads at the time of the crash. According to FARS manual, the longitude and latitude are based on the Police Crash Report, where the coordinates are either stated directly or geo-coded based on the accident address. We assume that the incidence of missing coordinates is caused by clerical or administrative omissions, unrelated to the combination of the victim's race and the hit-and-run status. For example, all the records from the state of Pennsylvania do not have coordinates.

crash, neighborhood, and pedestrian characteristics, and outcomes.

For each crash location we identify the corresponding census block group – the smallest geographical unit for which the U.S. Census Bureau publishes statistical data – adding the demographic characteristics from the 2010 Decennial Census and population income from the 2010 American Community Survey for the location of the crash. We also append the demographic (2010 Decennial Census) and income information (2016 American Community Survey) to the home ZIP Code of the driver.

Definition 2.1. Block group type:

- *White* block group, if the share of white residents is above 80%
- *Black* block group, if the share of Black residents is above 80%
- *White+Black* block group, if none of the above, but the combined share of Black and white residents is above 80%
- *Blacks-in-minority* block group, if none of the above and the share of Black residents is below 10%
- *Mixed* block group, if none of the above

The demographic characteristics of pedestrians and the spatial distribution of crashes naturally resemble the underlying population densities in the United States.[7] Unsurprisingly, crashes involving Black (white) pedestrians happen in the census block groups with a higher share of Black (white) residents (see Table 1, Panel A below). In Black block groups, most pedestrians who die are Black (85%), while in white block groups the share of Black pedestrian fatalities is just 8% (Table 2 below). Moreover, hit-and-runs are much more frequent in Black block groups (30% of total fatalities) than in white block groups (14%), while clearance rates are much lower in Black block groups (30% of hit-and-runs, compared to 55% in white block groups). Hence, on average, Black pedestrians are more likely to be victims of a hit-and-run (25% vs. 16% for white pedestrians) and their cases are less likely to be cleared.

7. Figure A.1 in Appendix A shows the geographical distribution of crashes for the contiguous United States. Additionally, we classify census block groups into five types, as defined below. Figure A.2 in Appendix A shows the map of fatal pedestrian crashes across the block group types in Cook County, Illinois, which encompasses Chicago.

Table 1: Descriptive statistics and covariate balance between white and Black pedestrian fatalities, unconditional and conditional on location and time fixed effects

	A. Descriptive statistics				B. Covariate balance after partial. out time and loc.(ℓ) FE			
	Mean values by P's race:				ℓ=contig.clust		ℓ=census tract	
	White	Black	diff	pval	diff	pval	diff	pval
Neighborhood (c. blgr):								
white population (shr)	0.66	0.35	0.318	0.000	0.053	0.000	0.007	0.000
black population (shr)	0.12	0.44	-0.328	0.000	-0.047	0.000	-0.007	0.000
urban area population (shr)	0.70	0.81	-0.105	0.000	-0.006	0.022	0.000	0.985
median hh inc. (th USD)	51.46	41.45	10.003	0.000	2.242	0.000	0.020	0.916
Crash characteristics:								
daylight (1/0)	0.23	0.17	0.060	0.000	0.004	0.251	-0.001	0.712
dark, lighted (1/0)	0.35	0.42	-0.076	0.000	-0.004	0.585	0.005	0.559
dark, no light (1/0)	0.37	0.36	0.016	0.053	0.001	0.937	-0.004	0.584
arterial road (1/0)	0.66	0.60	0.063	0.000	0.025	0.000	0.006	0.381
number of lanes	2.79	2.93	-0.137	0.000	-0.006	0.729	-0.013	0.524
not intersection (1/0)	0.74	0.75	-0.006	0.381	-0.004	0.533	-0.015	0.045
2-way road, no divider (1/0)	0.60	0.56	0.043	0.000	-0.001	0.943	-0.000	0.982
2-way road, divider (1/0)	0.35	0.39	-0.037	0.000	0.001	0.839	0.001	0.873
straight aligned road (1/0)	0.91	0.93	-0.014	0.002	0.001	0.888	0.000	0.995
flat road (1/0)	0.78	0.80	-0.018	0.009	0.001	0.832	0.001	0.831
crosswalk (1/0)	0.10	0.07	0.032	0.000	0.015	0.001	0.013	0.008
crossw. avail. unknown (1/0)	0.10	0.12	-0.016	0.003	-0.003	0.533	0.007	0.222
4-way intersection (1/0)	0.16	0.16	0.004	0.484	0.007	0.236	0.009	0.153
no traffic controls (1/0)	0.82	0.85	-0.025	0.000	-0.013	0.023	-0.015	0.021
rainy (1/0)	0.08	0.08	-0.001	0.749	0.000	0.917	-0.002	0.743
clear weather (1/0)	0.72	0.76	-0.037	0.000	-0.008	0.248	-0.004	0.588
road surface not dry (1/0)	0.16	0.14	0.015	0.012	0.006	0.324	0.003	0.683
Pedestrian characteristics:								
age (years)	48.90	43.26	5.643	0.000	3.039	0.000	1.340	0.000
senior citizen (1/0)	0.17	0.08	0.091	0.000	0.031	0.000	0.010	0.063
woman (1/0)	0.31	0.28	0.027	0.000	0.005	0.505	0.010	0.258
intoxicated (1/0)	0.25	0.22	0.027	0.000	0.016	0.009	0.015	0.052
died at scene, en route (1/0)	0.46	0.46	0.005	0.570	0.003	0.694	-0.000	0.973
Outcomes:								
hit-and-run (1/0)	0.16	0.25	-0.087	0.000				
D identified & h&r (1/0)	0.08	0.09	-0.010	0.032				
D identified (1/0), h&r only	0.48	0.35	0.129	0.000				
Number of observations	13,710	4,791						

Notes: Fatal one-driver-one-pedestrian crashes in the U.S. in 2010-2016; non-Hispanic white and Black pedestrians only; sources: FARS, U.S. Census. Panel A provides descriptive statistics using the raw FARS data. Panel B calculates differences in mean values after partialling out location and time fixed effects as defined in Section 3. Location fixed effects are at the level of contiguous clusters of census tracts (the first two columns) or at the level of census tracts (the remaining two columns). **Abbreviations**: P – pedestrian, D - driver, (1/0) – dummy variable, (shr) – share, hh – household, diff – the difference between the mean values for white and Black pedestrians, pval – p-value of the test H0: diff = 0.

The race of drivers and pedestrians are correlated: white pedestrians are more likely to be hit by drivers who live in white ZIP Codes, while the opposite holds true for Black pedestrians (Table 2). Crucially, hit-and-run cases involving Black pedestrians are much less likely to result in the driver being identified (36%), compared to cases involving white pedestrians (50%), as shown in Table 1, Panel A.

Table 2: Descriptive statistics for pedestrian fatalities: mean values by block group types

	Block group type:				
	White	B-in-Mi	Mixed	W+B	Black
N obs	6,113.0	4,521.0	3,029.0	3,525.0	1,281.0
Black ped. (%)	7.8	12.4	40.2	40.7	84.9
D's ZIP Code >80% white (%)	67.0	16.2	11.2	21.5	6.7
Median HH inc in blgr (th USD)	57.1	53.8	41.8	41.3	29.8
Hit-and-run (%)	13.7	18.9	21.5	18.3	29.7
D identified (% of hit-and-runs)	54.7	43.7	35.3	45.7	29.7

Sources: Fatal one-driver-one-pedestrian crashes in the U.S. in 2010-2016; non-Hispanic white and Black pedestrians only; sources: FARS, U.S. Census. **Abbreviations**: B-in-Mi – *blacks-in-minority*; , W+B – *white + Black*, D – driver, Median HH inc in blgr – the median household income per census block group, averaged across the block groups of the same type. Block group types are according to Definition 2.1.

Hit-and-runs are correlated with other risk factors for drivers: based on the sample of drivers that were eventually identified, those who flee are more likely to be younger, male, intoxicated, driving someone else's car, without a valid driver's license, speeding before the crash, or have records of previous violations (See Table A.1 in Appendix A).

Additionally, we use the CRIS database, administered by the Texas Department of Transportation, which collects the crash data submitted by local law enforcement officers in the State of Texas. It has a similar scope and structure to FARS, with a few notable differences. As discussed earlier, the CRIS data covers also non-fatal crashes. The race of pedestrians in CRIS is encoded from the police reports (as identified by the officer), as opposed to the death certificate used in FARS. Hence, the race of pedestrians may not necessarily coincide for the same cases in FARS and CRIS. Although methodological difference may introduce minor differences in the classification of pedestrians, these differences are not consequential for our separate analyses of the two datasets.

We restrict the CRIS data to a subset of vehicle-pedestrian crashes, which resulted in an incapacitating injury or death of the pedestrian (c. 30% of all injuries).[8] We focus only on incapacitating injuries because of the high stakes involved in the driver's decision to stay and the low risks of such accidents being under-reported to the police. For the period 2010-2016, there are in total $6,221$ cases in our restricted sample, out of which $5,387$ contain coordinates. Among those that contain coordinates, $1,766$ cases involve Black pedestrians with an average share of hit-and-runs at 25.5%, and $3,621$ cases involve white pedestrians with an average share of hit-and-runs at 16.4%. Roughly a third of the sample are fatalities.

Overall, the data points to substantial demographic and some behavioral differences between the drivers hitting white pedestrians and those hitting Black pedestrians. This implies that the empirical analysis must account for spatial racial segregation and other systematic differences in the characteristics of the crash and the victim, which is developed in the following section.

3 Empirical Setup

After a driver hits a pedestrian, they can either stay ($h\&r = 0$), or flee ($h\&r = 1$). The driver's decision to flee can be influenced by their own characteristics (d), location (ℓ), time (t), crash circumstances (z), race of the victim ($r \in \{B, W\}$), and other salient features of the victim (v) like gender and age. Let us denote all crash characteristics other than d and r as $x \equiv \{\ell, t, z, v\}$ and the two potential outcomes, given x, as $h\&r_x^{r=B}$ (if the victim is Black) and as $h\&r_x^{r=W}$ (if the victim is white), but we only observe $h\&r_x^{r=B}(r = B)$ and $h\&r_x^{r=W}(r = W)$. We are interested in the expected difference in the two potential outcomes.

From the descriptive evidence, we know that drivers from whiter neighborhoods tend to hit white pedestrians in white neighborhoods, and drivers from blacker neighborhoods tend to hit Black pedestrians in Black neighbor-

8. Incapacitating injuries include "severe laceration resulting in exposure of underlying tissues/muscle/organs or resulting in significant loss of blood; Broken or distorted extremity (arm or leg); Crush injuries; Suspected skull, chest, or abdominal injury other than bruises or minor lacerations; Significant burns (second and third degree burns over 10% or more of the body); Unconsciousness when taken from the crash scene; Paralysis" (from TxDOT's Instructions to police for reporting crashes, 2020 edition).

hoods. Hence, one cannot merely compare the prevalence of hit-and-runs for Black victims directly with white victims. However, once we fix the location, time, and other crash characteristics, we assume that the identity of the victim is independent from the identity of the driver. We can assume this because most traffic accidents are unintentional in nature, to the extent that drivers do not "choose" whom to hit, but rather happen to hit a person at random, among nearby pedestrians. In other words, we can compare the prevalence of hit-and-runs for Black and white victims who got into a crash under similar circumstances, in neighborhoods with similar characteristics, and comparable drivers-at-risk. Formally, our identification strategy relies on the "equal mixing" assumption:

$$\begin{aligned} & r \perp\!\!\!\perp d \mid x \\ & r \perp\!\!\!\perp h\&r_x^{r=B}, h\&r_x^{r=W} \mid x \end{aligned} \tag{1}$$

We estimate the following linear probability model using the sample of pedestrian fatalities:

$$\begin{aligned} h\&r_{\ell,g,t,i} = & \beta^{(hr)} I\{r = B\}_i + \vartheta\, ped\_char_i + \gamma\, crash\_char_i + \\ & + \eta\, blgr\_char_g + \tau_t + \alpha_\ell + \epsilon_{\ell,i} \end{aligned} \tag{2}$$

In Equation 2, $h\&r_{l,t,g,i}$ is an indicator variable for hit-and-run for pedestrian i in block group g at time t and location ℓ. The controls include *crash_char*: a vector of indicator variables for crashes on arterial roads, presence of traffic controls, marked crosswalks,road intersections, the number of lanes, as well as light conditions (daylight, dark+lighted, dark+not_lighted, dusk, dawn). The vector *ped_char* includes the pedestrian's salient characteristics: *sex*, *age* and *intoxication* status, as well as age^2, age^3 to capture differential effects on the driver's decision to flee or differences in mortality rates. The *blgr_char* includes census block group's demographic characteristics such as *%white*, *%Black*, *%white-Hispanic*, *%asian*, *%urban* residents, and the *Household median income.* The time fixed effects include 46 indicator variables for the *hour* × *weekend*, 6 indicator variables for the *day of the week*, 12 × 7 indicator variables for *month* × *year* to capture any monthly variations in the rate of hit-and-runs. As location fixed effects we use either (1) a contiguous racial cluster of census tracts (defined below), or (2) a census tract. Finally, $\epsilon_{\ell,i}$ is an error term clustered at the same level of the location fixed effect. Here, $\hat{\beta}^{(hr)}$ is the parameter of interest, which should capture the racial gap in hit-and-runs conditional on all the other characteristics x.

Defining location. We use two definitions of location to account for location fixed effects. The first one is based on clustering the census tracts in each county into contiguous areas, based on their demographic composition.[9] We use the clustering algorithm by Chodrow (2017b), which identifies clusters of spatially-contiguous, demographically-homogeneous communities (see Appendix C for additional details). This sort of demographic spatial clustering may be too loose in some cases, as it accounts for residential segregation but does not account, for example, for busy business areas, which may also change the underlying demographic of pedestrians and drivers at risk. Alternatively, we also use a narrower definition of location – census tracts as location fixed effects. The different definitions of location are shown for Cook County (IL) in Figure 1 (see a similar graph for New York County (NY) in Figure A.3 in Appendix A).

Table 1, Panel B, shows that controlling for location and time fixed effects makes nearly all crash characteristics to become balanced across pedestrian race, especially when using census tracts – the narrower definition of location. Hence, the looser definition of location (based on contiguous clusters of census tracts) is more likely to violate the "equal mixing" assumption. However, in densely-populated areas the narrower definition (census tracts) can represent too small geographical areas: their boundaries are defined to keep the population size within a census tract approximately constant. At the same time, some census tracts – containing busy business areas or dangerous roads – are naturally more likely to have several observations than other census tracts, creating questions about sample selection and external validity. Hence, we believe that we need to use both definitions: the looser one is more representative of the general population of traffic accidents, while the narrower one is more compliant with the equal mixing assumptions.

9. According to the U.S. Census Bureau, census tracts are "small, relatively permanent statistical subdivisions of a county or county equivalent and generally have a population size between 1,200 and 8,000 people, with an optimum size of 4,000 people. When first established, census tracts are to be as homogeneous as possible with respect to population characteristics, economic status, and living conditions" (US Census Bureau 1994).

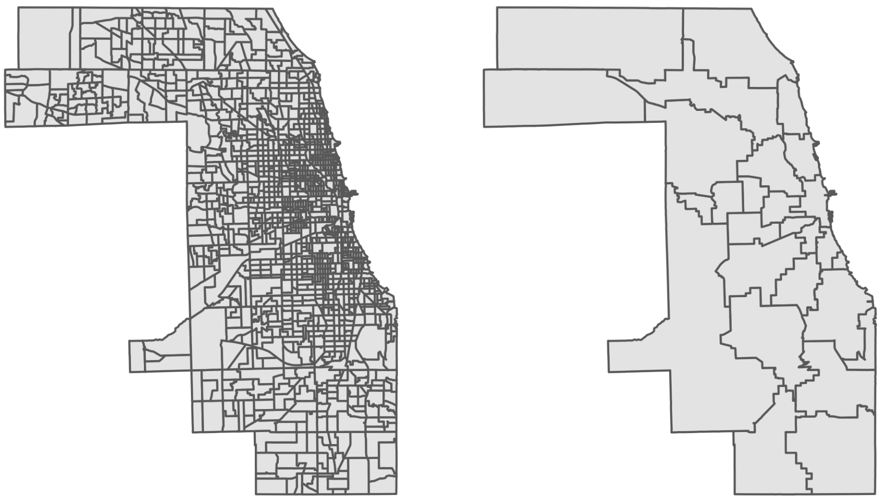

(a) Census tract

(b) Contiguous racial spatial cluster of census tracts

Figure 1: Two different definitions of location for Cook County (IL).

However, even when we use the narrow definition (census tract fixed effects), there is still statistical difference in the racial profile of the neighborhoods where the crash occur. In absolute terms, the difference is very small, < 1pp (e.g., compare to the difference of 32pp in the share of white residents in the raw data). Also, the location controls could not completely remove the difference in age: the Black population in the United States is simply younger than the white population. More importantly, there are still some minor differences in the intoxication status and the presence of crosswalks or traffic light controls – the variables that are probably more related to the behavior of pedestrians. Later in the results section, we show that the influence of these factors to the estimates of the racial gap in hit-and-run rates is very limited.

Selection into fatalities. Once we properly control for location, time, and other crash circumstances, $\beta^{(hr)}$ coefficient should capture the expected gap in hit-and-run rates for Black pedestrian *fatalities* relative to white pedestrian *fatalities*. We believe that the gap estimated by $\beta^{(hr)}$, based on fatalities, is a good approximation to the gap we would like to estimate, which includes all life-threatening crashes. In principle, it is not given that the two moments would coincide. For example, if the act of fleeing after hitting a pedestrian increased the chances that they died, then hit-and-run cases would be over-represented in the sample of fatal accidents: $\mathbb{E}(h\&r|r, x, died) \geq \mathbb{E}(h\&r|r, x)$. Moreover, the mortality rates may differ by race, if, for example, the extent of medical insurance coverage differs by race.[10] In this case, our estimator $\hat{\beta}^{(hr)}$ would be biased.

Nevertheless, we can show both theoretically (under certain general assumptions) and empirically (for the State of Texas) that our estimator $\beta^{(hr)}$ is, if anything, conservative, meaning that the sample of fatalities indeed can be used to test the behavior of drivers.

In particular, in Appendix B, we provide theoretical proofs of the following propositions. First, if drivers do not discriminate based on the pedestrian's race, then $\beta^{(hr)}$ should be equal to zero. If we reject the null hypothesis that $\beta^{(hr)} = 0$, then we can also reject the null hypothesis of no racial discrimination against Black pedestrians. Second, the sign of $\beta^{(hr)}$ provides information on the direction of racial discrimination: if $\beta^{(hr)} > 0$, drivers are more likely to flee after hitting Black pedestrians than white ones. Finally, the ratio of $\beta^{(hr)}$ to $\mathbb{E}(h\&r|r = W, x, died)$ provides a lower-bound estimate of the true discrimination rate. The reasoning is as follows: if hit-and-runs are more likely to cause death compared to hit-and-stays, then it will result in a higher share of hit-and-runs for white pedestrians conditional on fatality than unconditional on fatality, i.e., $\mathbb{E}(h\&r|r = W, x, died) \geq \mathbb{E}(h\&r|r = W, x)$. Notice that in the estimation of the relative racial gap,

$$\frac{\mathbb{E}(h\&r|r = W, x, died) - \mathbb{E}(h\&r|r = B, x, died)}{\mathbb{E}(h\&r|r = W, x, died)}$$

the baseline hit-and-run probability $\mathbb{E}(h\&r|r = W, x, died)$ enters both the denominator and numerator, and we can show that the partial derivative of the relative racial gap will always be negative with respect to the additional mortality caused by hit-and-runs. This conjecture assumes that hit-and-runs increase the mortality of pedestrians equally across races.

To illustrate, consider the following example. Assume that all hit-and-runs

10. Doyle Jr (2005) shows that in the U.S., the uninsured are more likely to die after a car crash than the insured, especially if the treatment calls for expensive procedures, like neurosurgery.

result in a death of the pedestrian, while only a half of hit-and-stays cause death. Also, assume that all drivers flee when the pedestrian is Black and only half of them flee when the pedestrian is white. Then, among fatalities, the share of hit-and-runs for Black pedestrians will be 100%, while the share of hit-and-runs for white pedestrians will be 2/3. Then the absolute gap in hit-and-run rates will be 1/3 based on fatalities, lower than the true absolute gap of 50%. Finally, the relative racial gap based on fatalities is 50% ($\frac{100\%-2/3}{2/3}$), which is lower than the true relative racial gap of 100% ($\frac{100\%-50\%}{50\%}$).

In the next section, we report the main results based on FARS data. Using the CRIS data for Texas, we show that indeed the point estimates of both the absolute and the relative racial gaps based on the fatal crashes are smaller in magnitude than the true estimates based on all life-threatening crashes, suggesting that, if anything, our estimates based on FARS data understates the full extent of racial gaps.

4 Results: racial gap in hit-and-run rates

Using FARS data, the results show that U.S. drivers are significantly more likely to flee after hitting a Black pedestrian, compared to a white pedestrian. The raw difference in the share of hit-and-runs between Black and white fatalities is 8.5 percentage points (see Column 1 of Table 3 below). As expected, comparing accidents that occur within the same contiguous cluster of census tracts and time fixed effects narrows down the raw differences in hit-and-runs: the racial gap drops to 3.2pp (see Column 2a). Adding controls for the differences in the characteristics of census block groups, crashes, and victims, as in Regression 2, further tightens the gap to 2.2 pp (Column 5a) with the p-value below 1%. In relative terms, it means that drivers are 12.8% more likely to flee after hitting a Black pedestrian than a white pedestrian. When we narrow the definition of location to census tracts, the point estimate jumps up to 3.5 percentage points and p-value to 2.5%, which puts the relative racial gap at 20.6% (Column 5b).

Table 3: Regression results: probability of a hit-and-run using different controls and definitions of location

	no controls	time, ℓ FE	+bl.gr. char.	+ped. char.	+crash char.	no ℓ FE
		A. ℓ = contig. cluster of tracts:				
	(1)	(2a)	(3a)	(4a)	(5a)	(6)
$\hat{\beta}^{(hr)}$	0.0853	0.0322	0.0262	0.0236	0.0222	0.0239
st. error	(0.0105)	(0.0082)	(0.0083)	(0.0083)	(0.0082)	(0.0078)
p-value	[0.0000]	[0.0001]	[0.0016]	[0.0045]	[0.0069]	[0.0020]
baseline*	0.1573	0.1710	0.1725	0.1732	0.1735	0.1731
relative gap	–	18.8%	15.2%	13.6%	12.8%	13.8%
N obs	18,043	18,043	18,043	18,043	18,043	18,043
N clusters	–	3,405	3,405	3,405	3,405	–
		B. ℓ = census tract:				
		(2b)	(3b)	(4b)	(5b)	
$\hat{\beta}^{(hr)}$		0.0402	0.0387	0.0350	0.0351	
st. error		(0.0158)	(0.0158)	(0.0158)	(0.0157)	
p-value		[0.0107]	[0.0142]	[0.0269]	[0.0251]	
baseline*		0.1689	0.1693	0.1702	0.1702	
relative gap		23.8%	22.9%	20.6%	20.6%	
N obs		18,043	18,043	18,043	18,043	
N clusters		13,947	13,947	13,947	13,947	
loc. (ℓ) FE	–	Yes	Yes	Yes	Yes	–
time FE	–	Yes	Yes	Yes	Yes	Yes
block gr. X	–	–	Yes	Yes	Yes	Yes
ped. X	–	–	–	Yes	Yes	Yes
crash X	–	–	–	–	Yes	Yes

* Predicted probability of hit-and-runs, if everyone in the data were white pedestrians: $\mathbb{E}(\widehat{h\&r}|r = W, x, died)$;
This table reports the results for regression 2. Column (1) provides estimates without any controls. Columns (2a)-(5a) report the results using the contiguous racial clusters of census tracts as the definition of location. Columns (2b)-(5b), instead, use census tracts as the definition of location. Standard errors – in parentheses – are clustered at the level of contiguous racial clusters of census tracts for columns (1), (2a)-(5a), and (6); and at the level of census tracts for columns (2b)-(5b). **Data**: Fatal one-driver-one-pedestrian crashes in the U.S. in 2010-2016; non-Hispanic white and Black pedestrians only; sources: FARS, U.S. Census.

This narrower definition of location is stricter in ensuring the "equal mixing" (conditional on location) assumption. Adding controls at the block group level to the contiguous clusters of census tracts reduces the estimated racial gap from 3.2pp to 2.6pp (see Columns 2a and 3a of Table 3), or approximately by 19%. By comparison, the same controls do not affect as much the estimates based on census tract fixed effects (compare Columns 3b to 2b). The latter is not surprising, since the narrower definition of location provides more balanced covariates, as shown earlier in Table 1, Panel B. Adding crash characteristics (e.g., road type, crosswalk) to other controls reduces the racial gap from 2.4pp to 2.2pp (or approximately by 6%, compare Columns 5a to 4a) when using the broader definition. These controls do not change the racial gap when using the narrower definition of location (compare Columns 5b to 4b). Therefore, the additional controls linked to the geographical differences and crash circumstances do not seriously affect our estimates based on the narrower definition of location. In light of this, we believe that the "equal mixing" assumption is justified, especially when defining location based on census tracts. Overall, adding extensive controls to census tract and time fixed effects only reduces the estimated magnitude of the racial gap by 13% (from 4.0pp to 3.5pp).

However, using census tracts comes at a cost in terms of sample size: only 14% of the sample actually contribute to the estimation of the racial gap in hit-and-runs (Table 3, Column 4). Out of about 18, 000 observations, only 2, 617 are from census tracts that have at least one white and one Black pedestrian fatality in the sample. These census tracts are more likely to be in urban areas, with a higher share of Black residents (see Table A.2 in Appendix A), which makes this estimate of the racial gap representative mostly of places with higher local mixing of Black and white pedestrians. Nevertheless, both definitions of location provide results that are in line with each other and not statistically different from each other.

The "equal mixing" assumption means that all characteristics of the pedestrian, besides race, should ex-ante be balanced across drivers. Hence, we can use the same regression setup to test whether drivers discriminate based on other salient characteristics of the pedestrian. Table A.3 in Appendix A shows coefficients for other pedestrian characteristics. We find that drivers do discriminate also based on age and intoxication status of the pedestrian. Drivers are significantly less likely to flee when hitting children and elderly: the point estimates for the polynomials of age suggest an inverted U-shape relationship between age and hit-and-runs, with a share of hit-and-runs around 6 percentage points higher for pedestrians in their thirties than 80-year-olds (or 5-year-olds). Interestingly, drivers are also less likely to flee when the pedestrian is intoxicated, perhaps because it is easier to attribute fault to the pedestrian in such cases, and the drivers may not fear legal repercussions. At the same time, we do not see that drivers discriminate based on gender. Hence, our test fails to reject the null hypothesis for at least one salient characteristic. The point estimate for the *female pedestrian* variable is small in absolute terms (at least when using contiguous clusters of census tracts as location definition) and in relative terms, compared to the estimated racial gap.

As previously stated, we control both for the location fixed effects and narrower block group characteristics, to mitigate any remaining spatial bias. However, controlling only for the block group characteristics also seem to be enough, as the results do not differ much when we include or exclude the fixed effects of the contiguous clusters (compare Columns 6 to 5a of Table 3).

The core specification, however, does not include interaction terms between the explanatory variables, which may result in an omitted-variable bias. By using the double-lasso method by Chernozhukov et al. (2015), we check whether our result (using contiguous clusters of census tracts as the definition of location) is robust to omitted interaction terms. As described in Appendix D, double lasso identifies 32 additional interaction terms that could predict either hit-and-runs or the race of the pedestrian. Adding them in Regression 2 does not change the result: it still puts the relative gap in victimization rate at 13.5% (See Column 2 of Table A.4 in Appendix A). We also estimate the racial gap nonparametrically using the generalized random forest (GRF) procedure (Athey et al. 2019). Since the GRF estimation is most useful for discovering heterogeneity in treatment effects, we describe this procedure in more detail in Section 5.1. Here, we want to point out that the overlap-weighted average treatment effect estimated by the GRF is similar to the OLS result: drivers are on average 14.8% more likely to flee when the pedestrian is Black rather than white (see Column 3 in Table A.4 in Appendix A). Finally, since our sample spans 7 years (a somewhat long time period), we split this time period in two or four subsamples, and then re-run the main regression using the interaction between location and time period fixed effects. Table A.5 show that our results do not change qualitatively when we restrict the time span.

Additionally, we study how the racial gap differs with light conditions at the time of the crash. We modify Regression 2, by interacting the race of the pedestrian with light conditions:

$$
\begin{aligned}
h\&r_{l,\ell,g,t,i} = \sum_{l} \beta_l^{(hr)} I\{r = B\}_i \times I(light = l) + \vartheta\, ped\_char_i + \\
+ \gamma\, crash\_char_i + \eta\, blgr\_char_g + \tau_t + \alpha_\ell + \epsilon_{\ell,i}
\end{aligned}
\tag{3}
$$

where l stands for the type of light condition at the time of the crash, in particular: daylight, dark (lighted), dark (no light), dawn, dusk, unknown lighting, other. Table A.6 in Appendix A shows the estimated coefficients for the three most common light conditions (daylight, dark with or without artificial lights), which together represent 95% of all accidents. The results show that the drivers discriminate more based on the pedestrian's race when there is artificial light, than when there is daylight or than when it is completely dark. However, the difference is not statistically significant once we partial out location fixed effects. Moreover, the difference in racial gaps across different light conditions most likely reflect differences in the population of drivers. For example, those who tend to crash on roads without lights might be systematically different than those on roads with lights. Also, it may reflect differences in the likelihood of having witnesses after the crash. For example, daylight crashes are more likely to have witnesses than night-time crashes, which may explain why racial gaps are higher at night. Unfortunately, we neither can account for differences in the population of drivers, nor does the data contains any information on the presence of by-standers.

To compare the estimates based on a selected sample of fatal cases with those based on all cases with incapacitating injuries, we also apply regression 2 to the CRIS dataset for Texas.[11] We find that in Texas, the relative racial gap is 24% using contiguous clusters fixed effects, and 31% using census tract fixed effects (See Table 4 below). Importantly, when we restrict the sample to fatal

11. Compared to our baseline specification, we omit the intoxication status of the pedestrian and the lane-width of the road from the regression, because of data limitations dataset. Also, instead of knowing whether the road is arterial or local, we only observe whether it belongs to the Texas Department of Transportation highway system or not.

cases, the point estimates of $\hat{\beta}^{(hr)}$ decrease slightly, but the standard errors also widen due to sample loss. As expected, the baseline share of hit-and-run cases is higher among fatalities than in the full sample, making the estimates of the relative racial gap based on fatalities more conservative than based on full data (at 21% and 24%, compared to 24% and 31%, respectively).

Table 4: Probability of a hit-and-run in Texas: all crashes with incapacitating injuries versus fatal crashes

	contig. cluster FE		census tract FE	
	all incap. injuries (1)	fatal cases (2)	all incap. injuries (3)	fatal cases (4)
$\hat{\beta}^{(hr)}$	0.0422	0.0386	0.0542	0.0438
standard error	(0.0130)	(0.0246)	(0.0190)	(0.0418)
p-value	[0.0012]	[0.1183]	[0.0044]	[0.2952]
baseline*	0.1782	0.1809	0.1737	0.1793
relative racial gap	23.7%	21.3%	31.2%	24.4%
N obs	5,230	1,757	5,274	1,780
N clusters	333	271	2,600	1,222
sample restricted to fatalities	–	Yes	–	Yes
controls[a]	Yes	Yes	Yes	Yes
location FE:				
cont. cluster of tracts FE	Yes	Yes	–	–
census tract FE	–	–	Yes	Yes

* Predicted probability of hit-and-runs, if everyone in the data were white pedestrians: $\mathbb{E}(\widehat{h\&r}|r = W, x)$;
[a] except for the intoxication status of the pedestrian
This table reports the results for regression 2 using the CRIS database. Columns (1) and (2) use the contiguous racial clusters of census tracts for location fixed effects, while Columns (3) and (4) use census tracts. Columns (1) and (3) use the sample of all cases that involve incapacitating injuries (fatal and non-fatal), while Columns (2) and (4) restrict the sample to fatal cases only. Standard errors – in parentheses – are clustered at the same level as the location definition. **Data**: One-driver-one-pedestrian crashes that involve incapacitating injuries to pedestrians, in the State of Texas in 2010-2016; white and Black pedestrians only; sources: CRIS database, U.S. Census.

Finally, we use CRIS to check whether the income of the pedestrian may explain some of the racial gap. So far we have controlled for such salient characteristics of the pedestrian as gender and age. However, the drivers may, in principle, act based on visual cues about the income level of the pedestrian. We proxy the income of the pedestrian with the per capita income in the ZIP Code of the pedestrian's home address. Table A.7 in Appendix A shows that adding this additional control does not change the results much. Taking the results at face value, the income of the pedestrian may explain only around 10% of the relative racial gap.

In the next subsection, we investigate whether clearance rates differ for white and Black hit-and-run victims.

4.1 Is there a racial gap in clearance rates?

After a hit-and-run, the driver is identified more often when the victim is white, rather than Black: the raw difference in clearance rates is 12.9 percentage points (Column 1 of Table 5 below). To test for a racial gap in clearance rates, we restrict the sample to hit-and-run cases only ($3,234$ observations) and we test whether the share of identified hit-and-run drivers is the same for Black and white victims, conditional on block group, crash, pedestrian characteristics, location, and time. We modify Regression 2 by changing the dependent variable to the indicator variable *identified*, which equals one if there is information about the hit-and-run driver.

$$\begin{aligned} identified_{\ell,g,t,i} =& \beta I\{r = B\}_i + \vartheta\, ped\_char_i + \gamma\, crash\_char_i + \\ & \eta\, blgr\_char_g + \alpha_\ell + \tau_t + \epsilon_{\ell,i} \end{aligned} \tag{4}$$

All other characteristics of the regression remain the same, except for the time controls: due to a smaller sample size we use *month* and *year* fixed effects separately, instead of *month*×*year* interaction terms. With this setup, parameter $\beta^{(id)}$ should capture the difference in the probabilities that the driver is identified when the victim is Black as opposed to when the victim is white, conditional on a fatality.

$$\beta^{(id)} = \mathbb{E}(identified | r = B; h\&r, x, died) - \mathbb{E}(identified | r = W; h\&r, x, died) \tag{5}$$

According to the results in Table 5, the gap in clearance rates shrinks to -5.4 percentage points (Column 5a) once we include the controls and the fixed effects of the contiguous racial clusters of census tracts. This estimate implies around 11.7% lower probability of identifying the hit-and-run offender for Black victims than for white victims. Again, the estimate is robust to different definitions of geographical locations. Although tightening the location definition increases the magnitude of the point estimate (to 6.8pp, Column 5b), it also substantially increases the standard errors, since the sample size is to small to use census tract fixed effects. Using contiguous clusters of census tracts, the estimates are only borderline significant (p-value of 6.7%) . When we add additional interaction terms selected by the double-lasso procedure, the regression returns an estimate of β at −5.5 percentage points, similar to the main specification (See Column 3 of Table A.9 in Appendix A). Generalized random forest estimation returns a slightly lower estimate of −4.4 percentage point and narrower standard errors, statistically significant at 5% level. Overall, the results suggest that the clearance rates are different for hit-and-run victims of different races.

As for the other characteristics of the pedestrian, we find that clearance rates are lower for intoxicated pedestrians and for elderly pedestrians (see Table A.8 in Appendix A). These are the groups of pedestrians who in the first place are less likely to become victims of hit-and-runs, conditional on being hit. Interestingly, the pedestrian's gender does not seem to play a role neither for hit-and-runs, nor for clearance rates.

One of the explanations to the racial gap in clearance rates could be the difference in the type of drivers who flee when the victim is Black rather than white, i.e., the presence of marginal drivers who react differently to the race of the victim, as we established in the previous section. Nevertheless, we find that the gap in the clearance rates is too large to be explained solely by the differences in the population of hit-and-run drivers. Consider the most conservative estimate in terms of magnitude produced by the GRF estimation (which is also the most precise) of a 9.6% lower clearance rate. Only if we assume that all of those marginal drivers – who contributed the additional 2.22 percentage points in hit-and-runs for Black pedestrians – are impossible to find, only then we can explain the gap in the clearance rate with just a simple explanation of the difference in the composition of drivers. In reality, it is unlikely that marginal hit-and-run drivers are much more effective in evading the police compared to average hit-and-run drivers (the baseline clearance rate is at 46% in Table 5).

In fact, when we look at how clearance rate is related to a gap in hit-and-run rates for the elderly or for intoxicated pedestrians, we may even suspect that the marginal hit-and-run drivers are in fact easier to find. Hence, some other channels, other than the inherent clearance rates for the marginal hit-and-run

Table 5: Regression results: probability that the driver is identified

	no controls	time, ℓ FE	+bl.gr. char.	+ped. char.	+crash char.	no ℓ FE
		A. ℓ = contig. cluster of tracts:				
	(1)	(2a)	(3a)	(4a)	(5a)	(6)
$\hat{\beta}^{(id)}$	-0.1316	-0.0479	-0.0424	-0.0530	-0.0540	-0.0439
st. error	(0.0203)	(0.0295)	(0.0297)	(0.0296)	(0.0295)	(0.0225)
p-value	[0.0000]	[0.1037]	[0.1533]	[0.0735]	[0.0673]	[0.0511]
baseline*	0.4898	0.4607	0.4588	0.4625	0.4628	0.4593
relative gap	–	-10.4%	-9.2%	-11.5%	-11.7%	-9.6%
N obs	3,234	3,234	3,234	3,234	3,234	3,234
N clusters		1,393	1,393	1,393	1,393	
		B. ℓ = census tract:				
		(2b)	(3b)	(4b)	(5b)	
$\hat{\beta}^{(id)}$		-0.0590	-0.0606	-0.0453	-0.0679	
st. error		(0.0736)	(0.0769)	(0.0795)	(0.0790)	
p-value		[0.4233]	[0.4305]	[0.5689]	[0.3903]	
baseline*		0.4645	0.4651	0.4598	0.4677	
relative gap		-12.7%	-13.0%	-9.9%	-14.5%	
N obs		3,234	3,234	3,234	3,234	
N clusters		3,023	3,023	3,023	3,023	
loc. (ℓ) FE	–	Yes	Yes	Yes	Yes	–
time FE	–	Yes	Yes	Yes	Yes	Yes
block gr. X	–	–	Yes	Yes	Yes	Yes
ped. X	–	–	–	Yes	Yes	Yes
crash X	–	–	–	–	Yes	Yes

* Predicted probability that the driver is identified, if everyone in the data were white pedestrians: $\mathbb{E}(\widehat{identified}|r = W, x, h\&r = 1, died)$;
This table reports the results for regression 4. Column (1) provides estimates without any controls. Columns (2a)-(5a) report the results using the contiguous racial clusters of census tracts as the definition of location. Columns (2b)-(5b), instead, use census tracts as the definition of location. Standard errors – in parentheses – are clustered at the level of contiguous racial clusters of census tracts for columns (1), (2a)-(5a), and (6); and at the level of census tracts for columns (2b)-(5b). **Data:** Fatal one-driver-one-pedestrian **hit-and-run** crashes in the U.S. in 2010-2016; non-Hispanic white and Black pedestrians only; sources: FARS, U.S. Census.

drivers, must be at play to explain the racial gap in clearance rates.

The racial gap in clearance rates may stem from a difference in the efforts and resources spent by the police, or from the efforts of the family of the deceased, or the cooperation of the local community in finding the offender, or from a combination of these factors. Unfortunately, our data is limited on what it can say on the causes of the clearance gap.

In summary, we find strong evidence that drivers tend to flee more often after hitting a Black pedestrian than after hitting a white pedestrian. We also find that there is a racial gap in the rate at which the hit-and-run cases are eventually cleared. The next section examines a few possible mechanisms to explain the relation between these findings.

5 Exploring potential mechanisms

This section analyzes two potential mechanisms behind the racial gap in victimization. The first, more direct one is *taste-based discrimination* (or "in/out-group bias"): under this mechanism, hit-and-runs are determined by the subjective preferences of drivers, who, following an accident, are more or less willing to extend assistance to their victims based on their race. The second, indirect mechanism is *statistical discrimination*: if cases involving Black victims are less likely to be solved – for whatever reason – drivers involved in accidents with Black pedestrians will expect a lower probability of being caught, and hence will be more likely to flee. We examine each mechanism in the next subsections, while Appendix E provides a simple Beckerian theoretical model to describe those mechanisms formally.

5.1 The racial gaps in hit-and-run rates by neighborhood type

If drivers discriminate by race because of preferences, then we should observe higher discrimination rates *against* Black pedestrians in places where drivers are predominantly non-Black – for example, on local roads in non-Black neighborhoods – and discrimination *in favor* of Black pedestrians in Black neighborhoods.

To estimate this, we expand the estimates of the moment of interest: whereas Regression 2 estimates an average effect $\beta^{(hr)}$, we now want to allow for the effect to vary with the characteristics of the crash, $\beta^{(hr)}(x)$. We use the generalized random forest (GRF) estimation (Athey et al. 2019), a non-parametric method that helps to account for potential heterogeneity in treatment effects and does not rely on functional form assumptions.[12] Helpfully, this method also points to the variables that best predict heterogeneous responses, which should reveal whether demographic characteristics of the neighborhoods are indeed predictive of differences in racial discrimination rates.

With the GRF, we estimate the following random effects model:

$$Y_i = b_i D_i + \epsilon_i; \quad \beta(x) = E(b_i | X_i = x); \tag{6}$$

where Y_i and D_i are the centered variables $Y_i = h\&r_i - \widehat{h\&r}(X_i)$ and $D_i = I\{r = B\}_i - I\widehat{\{r = B\}}(X_i)$, and $\widehat{h\&r}(X_i)$ and $I\widehat{\{r = B\}}_i$ are out-of-sample predictions using a random forest. Centering the variables $h\&r_i$ and $I\{r = B\}_i$ around their predicted values allows to remove the effect of confounders and concentrate on the heterogeneity in racial gaps.

In particular, the GRF estimates $\beta(x)$ non-parametrically:

$$\hat{\beta}(x) = \frac{\sum_{i=1}^{n} \alpha_i(x)(D_i - \bar{D}_\alpha)(Y_i - \bar{Y}_\alpha)}{\sum_{i=1}^{n} \alpha_i(x)(D_i - \bar{D}_\alpha)^2} \tag{7}$$

where $\alpha_i(x)$ are weights determined by the causal forest, and $\bar{D}_\alpha = \sum \alpha_i(x) D_i$ and $\bar{Y}_\alpha = \sum \alpha_i(x) Y_i$. We estimate Equation 7 by growing a causal forest consisting of five thousand trees.[13] The forest is cluster-robust with respect to the contiguous racial clusters: it takes clusters into account both at the bootstrapping stage and when estimating the variance.[14] We do not need to include interaction terms, as Random Forest automatically creates interaction terms while growing the trees. Instead of the contiguous racial clusters, we include

12. Instead of classical kernel weighting functions, GRF employs an adaptive weighting function based on a Random Forest algorithm. Random Forest is an ensemble learning method used for prediction (Breiman 2001). It aggregates information from many decision trees, each trained on random subsamples of the data, and also each time randomly restricting the set of variables evaluated for tree splits.

13. A causal forest grows many causal trees, each on a different random sample of observations. Each causal tree splits the sample into partitions (leaves) to maximize the estimated heterogeneity in treatment effects. The weight of observation i for a given $x \in X$ is determined by the frequency with which observation i falls into the same leaf as x, weighted across all the trees. See Athey et al. (2019) for more details on the weighting function and splitting rules.

14. We use the *GRF* package in R. See the guide in Tibshirani et al. (2020)

just the state identifiers and block group demographic characteristics, which earlier have shown to be sufficient controls for geographical differences in the underlying population of drivers and victims at risk.

The estimation is based on the overlap-weighted average treatment effect formula:[15]

$$\frac{\sum_{i=1}^{n}\left(e(X_i)(1-e(X_i))E[h\&r(r=B)-(h\&r(r=W))|X=X_i]\right)}{\sum_{i=1}^{n}\left(e(Xi)(1-e(Xi))\right)} \tag{8}$$

The GRF finds that the racial gaps in hit-and-runs vary the most by the demographic profile of the crash site. Since the causal forest for Regression 6 splits the data on those specific variables (and values) that maximize the difference in the racial gap, we look at the variables that were most often chosen by the forest to split the sample, as those variables are the most important for capturing heterogeneity in racial gaps. If we combine variables related to the demographic characteristics (percent of Hispanic, Asian, white, and Black population), household income, and urban residence, they account for 61% of all splits in the causal forest (See Table A.10 in Appendix A).[16]

Since the racial composition of neighborhoods in the U.S. is highly correlated with income, we further distinguish neighborhoods not just by racial composition – as defined in 2.1: *white, blacks-in-minority, mixed, white+Black, Black* – but also by income. We distinguish between low-income and high-income neighborhoods, namely areas with the median income above and below 49,445 USD (the national median for 2010). To more finely account for local geography, we also differentiate between arterial roads (i.e., high-capacity urban roads) and local roads (i.e., roads that allow access to property). The traffic on an arterial road may be quite different from the traffic on a local road, even if the two run close to each other, and we can expect that the race of the driver is more aligned with the racial profile of the neighborhood on local roads than on arterial roads.

By conditioning on the type of the neighborhood, the GRF reveals that the

15. We use the overlap-weighted formula since there are some areas for which the propensity score of hitting a Black pedestrian is quite close to zero or one, the overlap-weighted formula helps avoiding the division by the propensity score, as recommended by Li et al. (2018). where $e(x) = P[r_i = B|X_i = x]$

16. The other important variable is time of the accident (hour, light conditions, day of the week, month, year), which account for 22% of all splits.

strongest discrimination happens on local streets in poorer *white* neighborhoods, with the gap in hit-and-run rates between Black and white pedestrians at 17.4 percentage points (p-value < 1%). See Figure 2a below (and Table A.11 in Appendix A for more details). Given that the probability of a hit-and-run for white pedestrians in those neighborhoods is 13%, the results suggest that drivers on local roads in poorer white neighborhoods are more than twice as likely to flee when the pedestrian is Black rather than white. At the same time, there is no strong evidence of racial discrimination on local roads in richer *white* neighborhoods, but the point estimate still suggests a discrimination rate of around 20%. Similarly, the drivers on local roads in poorer *blacks-in-minority* and *mixed* neighborhoods also show a high racial discrimination against Black pedestrians, which is statistically significant. In those neighborhoods, the estimated relative gaps imply 40% higher probability of a hit-and-run for Black pedestrians in comparison to white pedestrians.

At the same time, there is no evidence that drivers racially discriminate against Black pedestrians on local streets in *Black* neighborhoods. If anything, the point estimates in *Black* neighborhoods and richer *mixed* neighborhoods suggest that the drivers might be discriminating against white pedestrians, however, the estimates are not statistically significant.

The GRF points to demographic characteristics of neighborhoods as the most likely drivers of heterogeneity in the racial gaps, but it does not establish that these differences are significant. To address this, we re-run the GRF estimation allowing the splits only on the type of the block group, income level (high/low), and the type of the road (arterial/local), and we formally test for heterogeneity using the calibration test by Chernozhukov et al. (2018).[17] Additional calibration test results (see Table A.12 in Appendix A) suggest that the random forest estimator has captured some heterogeneity in the drivers' racial discrimination of pedestrians, and that we can reject the null hypothesis of a homogeneous reaction across different neighborhoods at 5% significance level.

Could this racial gap in hit-and-run rates be driven by the fact that black pedestrians are less likely to be residents in non-black neighborhoods? While FARS does not contain data on the residence of the pedestrian, CRIS provides

17. The test by Chernozhukov et al. (2018) fits a linear model of the target estimand as a function of average racial discrimination – the mean GRF prediction – and the differential racial discrimination as estimated by the GRF. If the GRF captures no additional variation in racial discrimination, then the coefficient in front of the differential prediction will be zero.

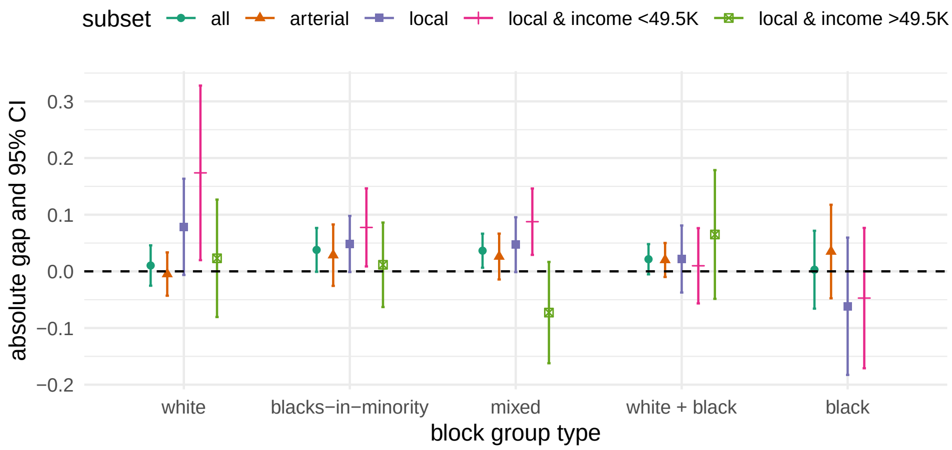

(a) Absolute racial gap in hit-and-run rates

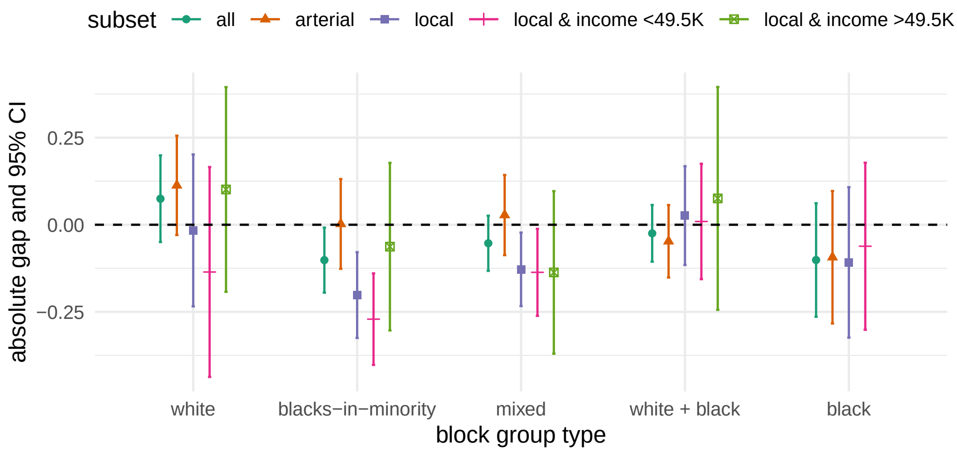

(b) Absolute racial gap in clearance rates

Figure 2: Racial gaps in hit-and-run and clearance rates by block group type
Estimates are based on the GRF estimation of Equation 7. See Table A.11 in Appendix A for more details. Data: Fatal one-driver-one-pedestrian crashes in the U.S. in 2010-2016; non-Hispanic white and Black pedestrians only; sources: FARS, U.S. Census.

the ZIP Code of the home address of the pedestrian in Texas. Augmenting Regression 2 by interacting the race of the pedestrian with their residence status, we find that the residence of the pedestrian does not explain the racial gap in hit-and-run rates in those neighborhoods (see Table A.13 in Appendix A).

In short, the racial gap in hit-and-run rates indeed varies across different types of neighborhood, along the lines of the in-group/out-group bias hypothesis. However, the observed heterogeneity does not prove *per se* that the gap is driven solely by the out-group bias of drivers. These differences could still be driven by statistical discrimination, if racial gaps in clearance rates also have the same pattern across neighborhoods. Indeed, this seems to be the case, as illustrated in the next subsection.

5.2 Statistical discrimination as a potential channel

We apply the GRF procedure to clearance rates (as opposed to hit-and-runs) and find that the racial gap in clearance rates is especially pronounced on local streets in *blacks-in-minority* and *mixed* neighborhoods (as shown in Figure 2 above and in Table A.14 in Appendix A). The gap in clearance rates in poorer white block groups is also salient, but not statistically significant and not as large in relative terms as *blacks-in-minority* and *mixed* block groups. The calibration test by Chernozhukov et al. (2018) tells us that this heterogeneity is marginally statistically significant (p-value of 5.1%, see Table A.15 in Appendix A). In general, the heterogeneity in clearance rate gaps across locations chimes with the heterogeneity in racial discrimination by drivers in hit-and-run decisions.

We ask whether the difference in clearance rates could be enough to explain the racial gap in hit-and-run rates. If statistical discrimination is at play, we should see higher racial gaps in hit-and-runs in those places with the biggest racial gap in clearance rates.

In general there is a negative relation between the out-of-sample predictions of hit-and-run rates ($\hat{y}^{hr}$) and that of clearance rates ($\hat{y}^{id}$), as shown in Figure 3a below. This is in line with the standard Beckerian prediction that crime rates will be higher when clearance rates are lower. Indeed, we also observe a negative relation in terms of *racial gaps* in hit-and-runs ($\beta^{(hr)}(x_i)$) and clearance rates ($\beta^{(id)}(x_i)$), as shown in Figure 3b. The negative slope is in line with the statistical discrimination hypothesis. Crucially, if the racial gap in hit-and-run rates was *exclusively* driven by statistical discrimination and drivers were fully-informed about the differences in clearance rates, then we would observe no racial gap in hit-and-runs when there is no racial gap in clearance rates, i.e., the distribution should be centered around the origin. This is not the case: even when the predicted racial gap in clearance rate is zero, the predicted racial gap in hit-and-run rates is expected to be positive.

We conclude that clearance rates may indeed explain some part of the racial

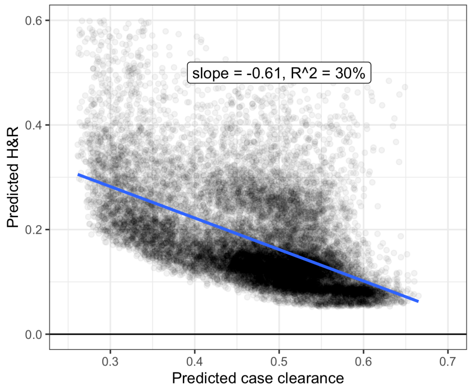

(a) Predicted clearance rate vs. predicted hit-and-run rates

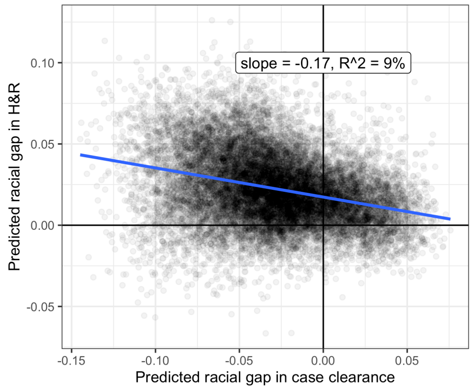

(b) Predicted racial gaps in clearance rates vs predicted racial gaps in hit-and-run rates

Sub-figure (a) uses the out-of-sample regression forest prediction of the hit-and-run rates and the prediction of the probability that the case is cleared, based on the characteristics of the case (excluding the race of the pedestrian). Sub-figure (b) uses the out-sample GRF prediction of the racial gap in hit-and-runs, $\hat{\beta}_i^{(hr)}$, and the racial gap in clearance rates, $\hat{\beta}_i^{(id)}$, for 18,043 cases. **Data:** Fatal one-driver-one-pedestrian hit-and-run crashes in the U.S. in 2010-2016; non-Hispanic white and Black pedestrians only; sources: FARS, U.S. Census.

discrimination by drivers, but there is still a considerable part of it that either stems from biased (inaccurate) beliefs about clearance rates, or taste-based discrimination (the evidence of which we find in Section 5.3), or some residual statistical discrimination due to other potential differences in the treatment of Black victims by law enforcement and criminal justice. The latter, however, seems less plausible, since growing body of research shows that the offenders of Black victims tend to receive milder punishment in court than the offenders of white victims (Glaeser and Sacerdote 2003; Alesina and La Ferrara 2014). In the next section, we test an out-group bias more directly.

5.3 Testing for out-group bias by looking at the composition of drivers

The results so far suggest that the strongest discrimination against Black pedestrians happens in poorer non-Black neighborhoods. However, comparing racial gaps in hit-and-run rates across locations does not give definitive answers on the presence of taste-based discrimination, since racial gaps in clearance rates also vary by location. Hence, we need to test whether non-Black drivers are more likely to discriminate against Black pedestrians than Black drivers are, within the same location.

We test for the presence of an out-group bias by comparing which drivers stay for white pedestrians and which drivers stay for Black pedestrians, within the same location. Under the "equal mixing" assumption, the race of the drivers should be independent of the race of the pedestrian, conditional on location, time, and crash characteristics. Hence, if drivers of any race discriminate against Black pedestrians to the same extent, then the racial profile of drivers who stay should be balanced across white and Black pedestrians. If, however, the decision to stay depends on the alignment between the race of the driver with the race of the pedestrian, then we should expect a higher share of white drivers among those who stay for white pedestrians, and of Black drivers for Black pedestrians.[18] In particular, in the CRIS dataset (for Texas), there is information on the race of the driver (as reported by the police officer), which we can use for our purposes.

We run the following regression on the sample of drivers who caused incapacitating injuries to pedestrians, but *stayed* (no hit-and-run):

$$\begin{aligned} D_{\ell,g,t,i} =& \beta^{(D)} I\{r = B\}_i + \vartheta \, ped\_char_i + \gamma \, crash\_char_i + \\ &+ \eta \, blgr\_char_g + \tau_t + \alpha_\ell + \epsilon_{\ell,i} \end{aligned} \tag{9}$$

where $D_{\ell,g,t,i}$ is an indicator variable for a certain race of the driver; All the other variables are defined the same as in Regression 2. The $\beta^{(D)}$ coefficient captures the expected gap in the share of drivers of race D for Black pedestrians relative to white pedestrians.

$$\beta^{(D)} = \mathbb{E}(D|r = B; x, h\&r = 0) - \mathbb{E}(D|r = W; x, h\&r = 0) \tag{10}$$

We restrict the sample for hit-and-stay cases where the driver is either white, Black, or Hispanic. Then, we run Regression 6 three times: once per each race of the driver as dependent variable. Table 6 below provides the results. If the higher gap in hit-and-runs was due to the expectations of racial gap in clearance rates in certain locations, then we would expect similar discrimination rates by all drivers independent of their own race. In other words, if there is no taste-based discrimination, then $\beta^{(D)}$ should be zero for everyone.

18. This test is similar to Kurmangaliyeva (2018).

Table 6: Regression results for the share of hit-and-stay drivers of certain race in Texas

	Dependent variable:		
Panel A: cont. cluster of tracts FE	White driver	Black driver	Hispanic driver
	(1)	(2)	(3)
$\hat{\beta}^{(D)}$	-0.0631	0.0638	-0.0007
standard error	(0.0176)	(0.0197)	(0.0149)
p-value	[0.0004]	[0.0013]	[0.9618]
baseline⋆	0.5685	0.2035	0.2280
N obs	4,017	4,017	4,017
N clusters	325	325	325
cont. cluster of tracts FE	yes	yes	yes
	Dependent variable:		
Panel B: census tract FE	White driver	Black driver	Hispanic driver
	(4)	(5)	(6)
$\hat{\beta}^{(D)}$	-0.0312	0.0436	-0.0124
standard error	(0.0273)	(0.0265)	(0.0241)
p-value	[0.2524]	[0.0998]	[0.6073]
baseline⋆	0.5589	0.2096	0.2315
N obs	4,017	4,017	4,017
N clusters	2,260	2,260	2,260
census tract FE	*yes*	*yes*	*yes*

⋆ Predicted probability that the driver is D, if everyone in the data were white pedestrians: $\mathbb{E}(\widehat{D}|r = W, x, h\&r = 0, died)$;
The table reports the estimate of β from Regression 9, where the dependent variable changes by column, and the independent variable of interest is the indicator variable for Black pedestrians. The regression includes core set of controls and uses the contiguous racial clusters of census tracts as location fixed effects (ℓ) for Panel A and census tracts as ℓ for Panel B. The sample is restricted to hit-and-stay cases only. Standard errors – in parentheses – are clustered at the level of location (ℓ). **Data:** One-driver-one-pedestrian crashes with incapacitating (fatal or not) injuries in Texas in 2010-2016; white, Hispanic, Black drivers only; non-Hispanic white and Black pedestrians only; sources: CRIS, U.S. Census.

Panel A of Table 6 suggests that white drivers are significantly under-represented among those who stay for Black pedestrians. Meanwhile, the share of Black drivers is significantly over-represented among those who stay for Black pedestrians. However, once we use tighter location fixed effects (See Panel B of Table 6) the point estimate of $\hat{\beta}^{(D)}$ for white drivers halves in magnitude from -6.3 percentage points to just -3.1pp and is no longer statistically different from zero. The point estimate of $\hat{\beta}^{(D)}$ for Black drivers also drops from 6.4pp to 4.3pp, and remains statistically significant only under 10% significance level.

Due to the importance of the "equal mixing" assumption, we prefer the more conservative test based on the narrower definition of location (Panel B), and interpret its results as only tentative evidence that the alignment in the race of the driver and the pedestrian is a factor in the driver's decision to stay or flee. While the signs of the coefficients are in line with the hypothesis of taste-based discrimination, the test fails to reject the null hypothesis at 5% significance level when using the narrow definition of location.

Overall, the results of Regression 9 corroborate the findings in Section 5.1, which showed that racial gaps in hit-and-runs vary most based on the racial composition of neighborhoods. Namely, Black pedestrians are discriminated more in poorer non-Black neighborhoods. However, independent of the location of the crash, non-Black drivers seem to be more likely to flee when the victim is Black, compared to Black drivers. The results mean that at least some part of the racial gap in hit-and-run rates might be attributed to out-group bias.

6 Conclusion

This paper provides the first causal test on whether decisions of offenders change with the race of the (potential) victim. Using data from Fatality Analysis Reporting System, we find that, all things being equal, drivers are 13% more likely to flee after hitting a Black pedestrian, compared to a white pedestrian. Moreover, the hit-and-run cases involving Black pedestrians are 12% less likely to be cleared, meaning the driver is not identified. The racial gap in hit-and-run rates is correlated with the racial gap in clearance rates across locations. This is especially the case on local roads in poorer non-Black neighborhoods. Moreover, we find tentative evidence that Black drivers are less likely to discriminate against Black pedestrians.

Our identification strategy requires that we observe all characteristics that are important for the decision of drivers to stay or flee. In our main regression, beside location and time, we control for the age, gender, jaywalking (i.e., presence of crosswalk), and intoxication status of the pedestrian, which we believe are salient at the moment of the crash. While the race of the pedestrian is a significant factor in the decision of drivers to flee, we show that gender is not.

To understand what drives the racial gap in hit-and-runs, we look at two potential types of discrimination that could be at play. We find suggestive evidence that, indeed, drivers might be statistically discriminating against Black pedestrians: the lower clearance rates for Black victims may induce some drivers to flee from Black pedestrians. This underlines the importance of setting a policy target to close the racial gap in clearance rates between Black and white victims. Hopefully, closing the gap in clearance rates may reduce the probability that drivers react differently based on the race of the pedestrian in the first place.

At the same time, statistical discrimination is likely not the only driver of the racial gap in hit-and-run rates, as taste-based-discrimination may also be at play. We find suggestive evidence of it in the fact that Black drivers are less likely to discriminate against Black pedestrians. In terms of policy prescriptions, statistical discrimination is arguably easier to tackle, by monitoring and attempting to close the racial gap in clearance rates.

Although racial discrimination has been documented in several stages of the American criminal-justice system, this paper provides the first evidence of racial gaps in victimization rates, namely discrimination by the general public, rather than institutional actors.

Our findings point to several areas of further research. More evidence is needed to establish to what extent discriminating behavior is caused by statistical or taste-based discrimination, and to understand the dynamics leading to the racial gap in clearance rates, along the lines of (Harvey and Mattia 2019). Lastly, we choose to look at hit-and-runs because they allow for plausibly exogenous matching between victims and offenders. However, racial gaps may well occur in other types of crimes, so more research is needed to establish the causal role of race in victimization in other settings and its dependence on clearance rates.

References

Alesina, A., and E. La Ferrara. 2014. "A test of racial bias in capital sentencing." *The American Economic Review* 104 (11): 3397–3433.

Anwar, S., and H. Fang. 2006. "An alternative test of racial prejudice in motor vehicle searches: Theory and evidence." *American Economic Review* 96 (1): 127–151.

Arrow, K., et al. 1973. "The theory of discrimination." *Discrimination in labor markets* 3 (10): 3–33.

Athey, S., J. Tibshirani, S. Wager, et al. 2019. "Generalized random forests." *The Annals of Statistics* 47 (2): 1148–1178.

Becker, G. 1971. *The economics of discrimination. chicago: Univ.*

Becker, G. S. 1974. "Crime and punishment : An economic approach." In *Essays in the Economics of Crime and Punishment,* Essays in, 1–54. UMI. ISBN: 0-87014-263-1.

Belloni, A., D. Chen, V. Chernozhukov, and C. Hansen. 2012. "Sparse models and methods for optimal instruments with an application to eminent domain." *Econometrica* 80 (6): 2369–2429.

Bertrand, M., and S. Mullainathan. 2004. "Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination." *American economic review* 94 (4): 991–1013.

Breiman, L. 2001. "Random forests." *Machine learning* 45 (1): 5–32.

Castriota, S., and M. Tonin. 2019. "Stay or Flee? Probability Versus Severity of Punishment in Hit-And-Run Accidents."

Chernozhukov, V., M. Demirer, E. Duflo, and I. Fernandez-Val. 2018. *Generic machine learning inference on heterogenous treatment effects in randomized experiments.* Technical report. National Bureau of Economic Research.

Chernozhukov, V., C. Hansen, and M. Spindler. 2016. "hdm: High-Dimensional Metrics." *R Journal* 8 (2): 185–199.

Chernozhukov, V., C. Hansen, and M. Spindler. 2015. "Post-selection and post-regularization inference in linear models with many controls and instruments." *American Economic Review* 105 (5): 486–90.

Chodrow, P. 2017a. *compx: Analyzing Spatial Segregation with Information Geometry.* https://philchodrow.github.io/compx/vignette_clustering.html. Accessed: 2020-04-08.

Chodrow, P. S. 2017b. "Structure and information in spatial segregation." *Proceedings of the National Academy of Sciences* 114 (44): 11591–11596.

Comino, S., G. Mastrobuoni, and A. Nicolò. 2020. "Silence of the innocents: Undocumented immigrants' underreporting of crime and their victimization." *Journal of Policy Analysis and Management* 39 (4): 1214–1245.

Cutler, D. M., E. L. Glaeser, and J. L. Vigdor. 1999. "The rise and decline of the American ghetto." *Journal of political economy* 107 (3): 455–506.

Doyle Jr, J. J. 2005. "Health insurance, treatment and outcomes: using auto accidents as health shocks." *Review of Economics and Statistics* 87 (2): 256–270.

Duarte, J., S. Siegel, and L. Young. 2015. "To lend or not to lend: revealed attitudes towards gender, ethnicity, weight, and age in the US."

Fagan, J., and A. Geller. 2018. "Police, Race, and the Production of Capital Homicides." *Berkeley J. Crim. L.* 23:261.

FARS: National Highway Traffic Safety Administration. Fatality Analysis Reporting System. https://www.nhtsa.gov/research-data/fatality-analysis-reporting-system-fars. Washington, DC: US Dept of Transportation. 2010-2016.

Feigenberg, B., and C. Miller. 2021. "Racial divisions and criminal justice: Evidence from southern state courts." *American Economic Journal: Economic Policy* 13 (2): 207–40.

Fong, C. M., and E. F. Luttmer. 2009. "What determines giving to Hurricane Katrina victims? Experimental evidence on racial group loyalty." *American Economic Journal: Applied Economics* 1 (2): 64–87.

French, M. T., and G. Gumus. 2015. "Hit-and-Run or Hit-and-Stay: Do Stricter BAC Limits Encourage Drivers to Flee the Crash Scene?" In *Proc., 2015 Fall Conference: The Golden Age of Evidence-Based Policy, Miami, FL.*

Glaeser, E. L., and B. Sacerdote. 2003. "Sentencing in homicide cases and the role of vengeance." *The Journal of Legal Studies* 32 (2): 363–382.

Grogger, J., and G. Ridgeway. 2006. "Testing for racial profiling in traffic stops from behind a veil of darkness." *Journal of the American Statistical Association* 101 (475): 878–887.

Guryan, J., and K. K. Charles. 2013. "Taste-based or statistical discrimination: the economics of discrimination returns to its roots." *The Economic Journal* 123 (572): F417–F432.

Harvey, A., and T. Mattia. 2019. "Reducing Racial Disparities in Crime Victimization."

Horrace, W. C., and S. M. Rohlin. 2016. "How dark is dark? Bright lights, big city, racial profiling." *Review of Economics and Statistics* 98 (2): 226–232.

Knowles, J., N. Persico, and P. Todd. 2001. "Racial bias in motor vehicle searches: Theory and evidence." *Journal of Political Economy* 109 (1): 203–229.

Kurmangaliyeva, M. 2018. "Missing Rich Offenders: Traffic Accidents and the Impartiality of Justice." *Review of Law & Economics* 14 (1).

Langley, M., and J. Sugarmann. 2017. "Black Homicide Victimization in the United States: An Analysis of 2014 Homicide Data."

Lee, C. 2005. "The value of life in death: Multiple regression and event history analyses of homicide clearance in Los Angeles County." *Journal of Criminal Justice* 33 (6): 527–534.

Levitt, S. D., and J. Porter. 2001. "How Dangerous Are Drinking Drivers?" *Journal of Political Economy* 109 (6): 1198–1237. eprint: https://doi.org/10.1086/323281.

Li, F., K. L. Morgan, and A. M. Zaslavsky. 2018. "Balancing covariates via propensity score weighting." *Journal of the American Statistical Association* 113 (521): 390–400.

Lowery, W., K. Kelly, T. Mellnik, and S. Rich. 2018. "Where killings go unsolved." *Washington Post* (June).

Manson, S., J. Schroeder, D. Van Riper, S. Ruggles, et al. 2017. "IPUMS National Historical Geographic Information System: Version 12.0 [Database]." *Minneapolis: University of Minnesota* 39.

Persico, N., and P. Todd. 2006. "Generalising the hit rates test for racial bias in law enforcement, with an application to vehicle searches in Wichita." *The Economic Journal* 116 (515): F351–F367.

Shayo, M., and A. Zussman. 2011. "Judicial ingroup bias in the shadow of terrorism." *Quarterly Journal of Economics* (126): 1447–1484.

Tibshirani, J., S. Athey, and S. Wager. 2020. *grf: Generalized Random Forests.* R package version 1.2.0.

Tibshirani, R. 1996. "Regression shrinkage and selection via the lasso." *Journal of the Royal Statistical Society: Series B (Methodological)* 58 (1): 267–288.

US Census Bureau. 1994. *Geographic areas reference manual.*

West, J. 2015. "Racial bias in police investigations." *Unpublished manuscript.*

A (Online Appendix) Additional graphs and tables

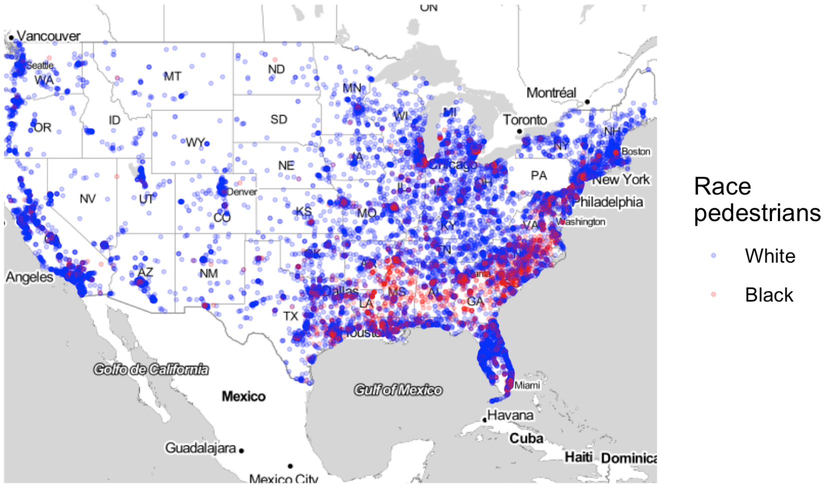

Figure A.1: Location of fatal traffic accidents reported in FARS across 48 contiguous U.S. states, 2010–2016. Source: FARS, shapefiles from NHGIS

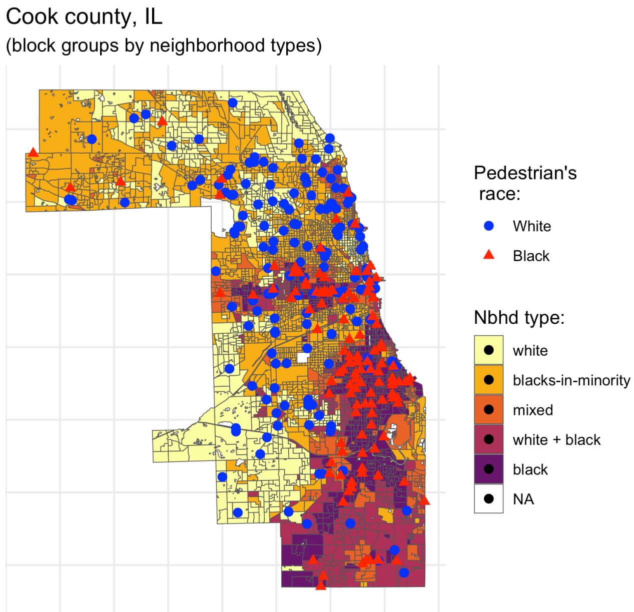

Figure A.2: Cook county, Illinois: block group types and pedestrian fatalities by race.
Sources: FARS, U.S. Census, NHGIS.

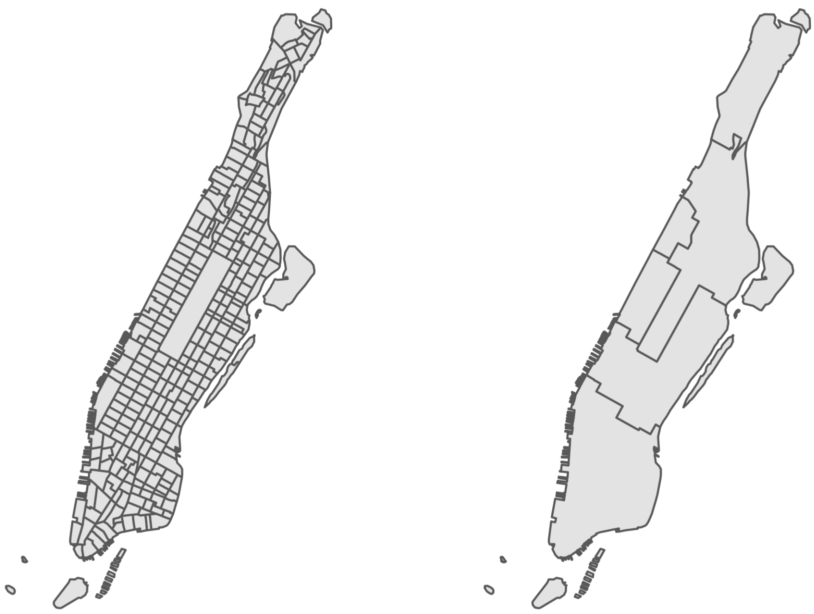

(a) Census tract

(b) Contiguous racial spatial cluster of census tracts.

Figure A.3: Two different definitions of location for New York County (NY).

Table A.1: Descriptive statistics for drivers involved in crashes with pedestrian fatalities: mean values by pedestrian's race

	White P	Black P	diff	pval
Drivers who stayed:				
age (years)	42.85	41.09	1.763	0.000
woman (1/0)	0.30	0.30	0.003	0.701
D.U.I. (1/0)	0.07	0.07	0.003	0.546
owner of the car (1/0)	0.62	0.58	0.034	0.000
business car (1/0)	0.12	0.11	0.005	0.394
valid driver license (1/0)	0.94	0.90	0.041	0.000
previous records (1/0)	0.41	0.44	-0.031	0.001
speeding before crash (1/0)	0.04	0.03	0.004	0.258
home ZC's white pop (shr)	0.67	0.47	0.198	0.000
home ZC's Black pop (shr)	0.12	0.32	-0.207	0.000
income per capita (th USD)	28.51	25.11	3.394	0.000
Drivers who run away, but were identified:				
age (years)	36.00	35.17	0.832	0.316
woman (1/0)	0.24	0.23	0.013	0.587
D.U.I. (1/0)	0.31	0.25	0.056	0.027
owner of the car (1/0)	0.51	0.43	0.079	0.006
business car (1/0)	0.06	0.06	-0.006	0.639
valid driver license (1/0)	0.73	0.67	0.067	0.012
previous records (1/0)	0.53	0.53	0.008	0.783
speeding before crash (1/0)	0.14	0.15	-0.007	0.728
home ZC's white pop (shr)	0.64	0.45	0.197	0.000
home ZC's Black pop (shr)	0.13	0.34	-0.216	0.000
income per capita (th USD)	26.86	24.32	2.548	0.000

Notes: Fatal one-driver-one-pedestrian crashes in the U.S. in 2010-2016; non-Hispanic white and Black pedestrians only; sources: FARS, U.S. Census. **Abbreviations**: P – pedestrian, (1/0) – dummy variable, (shr) – share, D.U.I. – driving under influence, ZC – ZIP Code.

Table A.2: Descriptive statistics for census tracts that have both white and Black pedestrian fatalities

	Census tracts:	
	At least one Black and one white P fatality	The rest
	mean	mean
urban area population in blgr (shr)	0.81	0.86
white population in blgr (shr)	0.60	0.45
Black population in blgr (shr)	0.18	0.30
Hispanic population in blgr (shr)	0.08	0.09
median hh income in blgr (th USD)	49.77	43.42
arterial road (1/0)	0.63	0.70

Data: Fatal one-driver-one-pedestrian crashes in the U.S. in 2010-2016; non-Hispanic white and Black pedestrians only; sources: FARS, U.S. Census. **Abbreviations**: (1/0) – dummy variable, (shr) – share, hh – household, blgr – census block group, P - pedestrian.

Table A.3: Probability of a hit-and-run: the linear probability model's coefficients for all pedestrian characteristics

	A. ℓ = contig. cluster			B. ℓ = census tract		
	coef.	se	p-val	coef.	se	p-val
Black pedestrian	0.0222	(0.0082)	[0.0069]	0.0351	(0.0157)	[0.0251]
female pedestrian	-0.0016	(0.0065)	[0.8050]	-0.0187	(0.0128)	[0.1436]
Pedestrian's age:						
age	0.0052	(0.0017)	[0.0021]	0.0033	(0.0036)	[0.3560]
age^2	-1.1e-04	(3.8e-05)	[0.0058]	-6.8e-05	(7.9e-05)	[0.3912]
age^3	5.4e-07	(2.6e-07)	[0.0338]	3.1e-07	(5.3e-07)	[0.5501]
H0: $age=age^2=age^3=0$			[0.0000]			[0.0400]
intoxicated ped.	-0.0404	(0.0081)	[0.0000]	-0.0389	(0.0138)	[0.0048]
main controls	Yes			Yes		
time FE	Yes			Yes		
location FE:						
cont. cluster	Yes			–		
census tract	–			Yes		

Similar to Table 3, this table reports results for regression 2, but it also reports coefficients for gender, age, and intoxication status of the pedestrian. Panel A uses the contiguous racial clusters of census tracts for location fixed effects, while Panel B uses census tracts. Standard errors – in parentheses – are clustered by location ℓ. P-values (H0: coef.=0) are in brackets. **Data**: Fatal one-driver-one-pedestrian crashes in the U.S. in 2010-2016; non-Hispanic white and Black pedestrians only; sources: FARS, U.S. Census.

Table A.4: Hit-and-runs: regression results with different specifications

	core LPM (1)	double-lasso LPM (2)	GRF (3)
$\hat{\mathbb{E}}\left[\hat{\beta}^{(hr)}\right]$: racial gap in hit-and-runs	0.0222	0.0233	0.0256^{a}
standard error	(0.0082)	(0.0082)	(0.0078)
p-value	[0.0069]	[0.0050]	[0.0011]
baseline*	0.1735	0.1730	0.1727
relative gap	12.8%	13.5%	14.8%
N obs	18,043	18,043	18,043
N clusters	3,405	3,405	3,405
main controls	Yes	Yes	Yes
interaction terms	–	Yes	Yes^{b}
state fixed effects	–	–	Yes^{c}
contiguous cluster of tracts FE	Yes	Yes	$–^{d}$

Column (1) reports the estimate of β from Regression 2 with the core set of controls and using the contiguous racial clusters of census tracts as location fixed effects (ℓ). Column (2) reports the estimate of β from Regression 2 with the additional interaction terms which have been selected by the double-lasso selection procedure for causal inference (Chernozhukov et al. 2015) (For the full set of additional interaction terms see Footnotes D and D). Column (3) reports the overlap-weighted average treatment effect based on the generalized random forest procedure (Athey et al. 2019), which estimates the random effects model using Equation 7 and weights α_i automatically determined by a random forest.

* Predicted probability of a hit-and-run, if everyone in the data were white pedestrians: $\mathbb{E}(\widehat{h\&r}|r = W, x, died)$;

[a] overlap-weighted average effect of the pedestrian's race, according to Equation 8

[b] the interaction terms are accounted for both at the prediction stage of the grf (thanks to double-lasso procedure) and at the causal-tree stage of the grf, as random forest create automatically interaction terms within the routine of individual decision trees.

[c] State identifiers are included only in the causal forest part of the grf procedure (to capture heterogeneity across states).

[d] uses the cluster-robust forests, which account for the contiguous racial clusters in the bootstrapping procedures and in variance estimation. Standard errors – in parentheses – are clustered at the same level of the location fixed effect.

Data: Fatal one-driver-one-pedestrian crashes in the U.S. in 2010-2016; non-Hispanic white and Black pedestrians only; sources: FARS, U.S. Census.

Abbreviations: LPM – linear probability model; *relative gap* $\equiv\mathbb{E}(\hat{\beta}_i)/\,\mathbb{E}(h\&r|r = W)$; $h\&r$ – indicator for a hit-and-run; $r = B$ – the pedestrian is a Black person; $r = W$ – the pedestrian is a non-Hispanic white person

Table A.5: Hit-and-runs: regression results with location×time-period fixed effects

	ℓ = contig. cluster of c. tracts			ℓ = census tract		
FE:	ℓ	$\ell \times 1/2T$	$\ell \times 1/4T$	ℓ	$\ell \times 1/2T$	$\ell \times 1/4T$
	(1)	(2)	(3)	(4)	(5)	(6)
$\hat{\beta}^{(hr)}$	0.0222	0.0200	0.0168	0.0351	0.0518	0.0430
standard error	(0.0082)	(0.0088)	(0.0098)	(0.0157)	(0.0202)	(0.0288)
p-value	[0.0069]	[0.0233]	[0.0864]	[0.0251]	[0.0104]	[0.1355]
baseline*	0.1735	0.1741	0.1749	0.1702	0.1659	0.1682
relative gap	12.8%	11.5%	9.6%	20.6%	31.2%	25.6%
N obs	18,043	18,043	18,043	18,043	18,043	18,043
N clusters	3,405	5,182	7,335	13,947	15,604	16,664
main controls	Yes	Yes	Yes	Yes	Yes	Yes
time FE	Yes	Yes	Yes	Yes	Yes	Yes
location FE:						
contig. cluster	Yes	–	–	–	–	–
contig. cluster × 1/2-period	–	Yes	–	–	–	–
contig. cluster × 1/4-period	–	–	Yes	–	–	–
census tract	–	–	–	Yes	–	–
census tract × 1/2-period	–	–	–	–	Yes	–
census tract × 1/4-period	–	–	–	–	–	Yes

* Predicted probability of a hit-and-run, if everyone in the data were white pedestrians: $\mathbb{E}(\widehat{h\&r}|r = W, x, died)$; Columns (1) and (4) report the estimate of β from Regression 2 with the core set of controls and location and time fixed effects. Columns (2) and (5) add an interaction term between the location and the time period split in two halves (time period 1 includes accidents that happened between 2010 and 2013, and time period 2 includes accidents that happened between 2014 and 2016). Columns (3) and (6), instead, splits the time period into four quarters. Columns (1)-(3) use contiguous clusters of census tracts as the definition of location, while Columns (4)-(5) use census tracts.

Table A.6: The racial gap in hit-and-run rates at different light conditions

		location fixed effects ℓ:		
	no controls (1)	contig. clusters (2)	census tracts (3)	no loc. FE (4)
E1: $\hat{\beta}^{(hr)}_{\text{daylight}}$	0.0619	0.0158	0.0221	0.0155
	(0.0140)	(0.0158)	(0.0349)	(0.0142)
	[0.0000]	[0.3173]	[0.5256]	[0.2762]
relative gap: daylight		10.3%	14.5%	10.5%
E2: $\hat{\beta}^{(hr)}_{\text{dark+light}}$	0.1059	0.0327	0.0407	0.0474
	(0.0147)	(0.0125)	(0.0232)	(0.0117)
	[0.0000]	[0.0089]	[0.0794]	[0.0001]
relative gap: dark, lighted		17.0%	21.9%	24.8%
E3: $\hat{\beta}^{(hr)}_{\text{dark+no light}}$	0.0499	0.0116	0.0365	0.0011
	(0.0127)	(0.0130)	(0.0221)	(0.0117)
	[0.0001]	[0.3741]	[0.0977]	[0.9253]
relative gap: dark, no light		7.2%	22.8%	24.8%
p-value H0: E1=E2=E3	[0.0030]	[0.4647]	[0.8926]	[0.0125]
N obs	18,043	18,043	18,043	18,043
N clusters	–	3,405	13,947	–
main controls	–	Yes	Yes	Yes
time FE	–	Yes	Yes	Yes
location FE:				
cont. cluster of tracts FE	–	Yes	–	–
census tract FE	–	–	Yes	–

This table reports the estimates for race×light coefficients for Regression 3. Column (1) provides estimates without any controls. Column (2) reports the results with the full set of controls and uses the contiguous racial clusters of census tracts for location fixed effects, while Column (3) uses census tracts. Column (4) does not use location fixed effects. Standard errors – in parentheses – are clustered at the level of contiguous racial clusters of census tracts for columns (1), (2), and (4), and at the level of census tracts for column (3). P-values (H0: coef.=0) are in brackets. **Data**: Fatal one-driver-one-pedestrian crashes in the U.S. in 2010-2016; non-Hispanic white and Black pedestrians only; sources: FARS, U.S. Census.

Table A.7: Probability of a hit-and-run in Texas: including the average income of the pedestrian's ZIP Code

	contig. clusters FE		census tract FE	
	core specif.	with P's ZIP Code income	core specif.	with P's ZIP Code income
	(1)	(2)	(3)	(4)
$\hat{\beta}$	0.0364	0.0320	0.0437	0.0383
standard error	(0.0144)	(0.0148)	(0.0215)	(0.0217)
p-value	[0.0117]	[0.0313]	[0.0419]	[0.0772]
baseline⋆	0.1695	0.1709	0.1667	0.1684
relative racial gap	21.5%	18.7%	26.2%	22.7%
N obs	4,356	4,356	4,396	4,396
N clusters	328	328	2,385	2,385
income in P's ZIP Code	–	Yes	–	Yes
main controls	Yes	Yes	Yes	Yes
location FE:				
cont. cluster of tracts FE	Yes	Yes	–	–
census tract FE	–	–	Yes	Yes

⋆ Predicted probability of a hit-and-run, if everyone in the data were white pedestrians: $\mathbb{E}(\widehat{h\&r}|r = W, x)$;
This table reports the results for regression 2 using the CRIS database. In all columns the sample is restricted to $4,356$ observations that contain ZIP Code information for the pedestrian with incapacitating injuries, including fatalities. Columns (1) and (2) use the contiguous racial clusters of census tracts for location fixed effects, while Columns (3) and (4) use census tracts. Columns (1) and (3) use the core regression specification, while Columns (2) and (4) add a control for the per capita income in the ZIP Code area of the pedestrian's residence. Standard errors – in parentheses – are clustered at the same level as the location definition.
Data: One-driver-one-pedestrian crashes that involve incapacitating injuries to pedestrians, in the State of Texas in 2010-2016; white and Black pedestrians only; sources: CRIS database, U.S. Census.

Table A.8: Probability that the driver is identified: linear probability model's coefficients for all pedestrian characteristics

	ℓ = contig. cluster			ℓ = census tract		
	coef.	se	p-val	coef.	se	p-val
Black pedestrian	-0.0540	(0.0295)	[0.0673]	-0.0679	(0.0790)	[0.3903]
female pedestrian	0.0131	(0.0231)	[0.5710]	-0.0176	(0.0641)	[0.7842]
Pedestrian's age:						
age	-0.0043	(0.0078)	[0.5812]	0.0018	(0.0250)	[0.9411]
age^2	1.1e-04	(1.7e-04)	[0.5455]	1.3e-04	(5.3e-04)	[0.8090]
age^3	-1.1e-06	(1.2e-06)	[0.3826]	-2.7e-06	(3.5e-06)	[0.4383]
H0: $age=age^2=age^3=0$			[0.0059]			[0.0160]
intoxicated ped.	-0.0780	(0.0286)	[0.0065]	-0.0050	(0.0788)	[0.9496]

Similar to Table 5, this table reports results for regression 4, but it also reports coefficients for gender, age, and intoxication status of the pedestrian. Column (1) provides estimates without any controls. Column (2) reports the results with the full set of controls. Column (3) uses the contiguous racial clusters of census tracts for location fixed effects, while Column (4) uses census tracts. Standard errors – in parentheses – are clustered at the same level of the location fixed effect, for columns (1) and (2) at the level of contiguous racial clusters of census tracts. P-values (H0: coef.=0) are in brackets. **Data**: Fatal one-driver-one-pedestrian **hit-and-run** crashes in the U.S. in 2010-2016; non-Hispanic white and Black pedestrians only; sources: FARS, U.S. Census.

Table A.9: Probability that the driver is identified: regression results with different specifications

	Core LPM (1)	Double-lasso LPM (2)	GRF (3)
$\mathbb{E}\left[\hat{\beta}^{(id)}\right]$: racial gap in clearance rates	-0.0540	-0.0548	-0.0448[a]
standard error	(0.0295)	(0.0297)	(0.0228)
p-value	[0.0673]	[0.0652]	[0.0490]
baseline[⋆]	0.4628	0.4631	0.4650
relative gap	-11.7%	−11.8%	−9.6%
N obs	3,234	3,234	3,234
N clusters	1,393	1,393	1,393
main controls	Yes	Yes	Yes
interaction terms	–	Yes	Yes[b]
state fixed effects	–	–	Yes[c]
contiguous cluster of tracts FE	Yes	Yes	–[d]

Column (1) reports the estimate of β from Regression 2 with the core set of controls and using the contiguous racial clusters of census tracts as location fixed effects (g). Column (2) reports the estimate of β from Regression 2 with the additional interaction terms which have been selected by the double-lasso selection procedure for causal inference (Chernozhukov et al. 2015) (For the full set of additional interaction terms see Footnotes D and D). Column (3) reports the overlap-weighted average treatment effect based on the generalized random forest procedure (Athey et al. 2019), which estimates the random effects model using Equation 7 and weights α_i automatically determined by a random forest.

⋆ Predicted probability that the driver is identified, if everyone in the data were white pedestrians: $\mathbb{E}(\widehat{identified}|r = W, x, h\&r = 1, died)$;

a overlap-weighted average effect of the pedestrian's race, according to to Equation 8.

b the interaction terms are accounted for both at the prediction stage of the grf (thanks to double-lasso procedure) and at the causal-tree stage of the grf, as random forest automatically creates interaction terms within the routine of individual decision trees.

c State identifiers are included only in the causal forest part of the grf procedure (to capture heterogeneity across states).

d uses the cluster-robust forests, which account for the contiguous racial clusters in the bootstrapping procedures and in variance estimation. Standard errors – in parentheses – are clustered at the same level of the location fixed effect.

Data: Fatal one-driver-one-pedestrian crashes in the U.S. in 2010-2016; non-Hispanic white and Black pedestrians only; sources: FARS, U.S. Census.

Abbreviations: LPM – linear probability model; *relative gap* $\equiv\mathbb{E}(\hat{\beta}_i)/\mathbb{E}(h\&r|r = W)$; $h\&r$ – indicator for a hit-and-run; $r = B$ – the pedestrian is a Black person; $r = W$ – the pedestrian is a non-Hispanic white person

Table A.10: Exploring heterogeneity of the racial gap in hit-and-runs with Generalized Random Forest: variables that correlate the most with the racial gaps in hit-and-run rates

	Variable[a]	Share of splits[b]	Most freq. split value[c]	Effect of P's race on h&r when ≤ S		when > S	
				mean	se	mean	se
1	Hisp. population (shr blgr)	0.17	0.283	0.013	0.008	0.073	0.035
2	Asian population (shr blgr)	0.12	0.001	0.046	0.035	0.014	0.008
3	white population (shr blgr)	0.11	0.441	0.039	0.013	0.004	0.011
4	median hh inc (th USD, blgr)	0.10	18.194	0.059	0.039	0.015	0.008
5	hour (0 to 23)	0.09	22	0.015	0.008	0.043	0.039
6	Black population (shr blgr)	0.09	0.656	0.020	0.009	-0.013	0.027
7	month (1 to 12)	0.05	11	0.010	0.009	0.077	0.027
8	pedestrian's age	0.05	20	0.007	0.024	0.018	0.009
9	day of week (1 to 7; 6 = Friday)	0.04	6	0.016	0.009	0.023	0.021
10	urban population, (shr blgr)	0.02	0.997	-0.006	0.015	0.029	0.010
11	year	0.02	2015	0.019	0.009	0.005	0.020
12	Texas state (1/0)	0.02	0	0.014	0.008	0.043	0.026
13	number of lanes	0.02	3	0.017	0.010	0.019	0.014
14	dark, lighted (1/0)	0.01	0	0.006	0.011	0.033	0.013
15	arterial road (1/0)	0.01	0	0.033	0.013	0.008	0.010
16	dark, no light (1/0)	0.01	0	0.030	0.010	-0.002	0.014

[a] Only the variables with the share of splits greater than 1%;

[b] A weighted sum of how many times variable *j* was split on at each depth in the causal forest estimation of Regression 6;

[c] The mode of split values for a given variable using the data from the very first split of each of the 5000 causal trees of the causal forest for Regression 6.

Data: Fatal one-driver-one-pedestrian crashes in the U.S. in 2010-2016; non-Hispanic white and Black pedestrians only; sources: FARS, U.S. Census. **Abbreviations**: (1/0) –- dummy variable, (shr) –- share, hh inc –- household income, blgr – census block group, P – pedestrian.

Table A.11: Exploring heterogeneity of the racial gap in hit-and-runs with Generalized Random Forest: estimates by the type of block group and road.

	Blgr type	Sample subset	Estimate	S.E.	p-value	Baseline	Rel. gap	n obs
1	white	all	0.010	0.018	0.57	0.13	0.08	6009
2	white	arterial	-0.005	0.019	0.81	0.13	-0.04	4529
3	white	local	0.079	0.043	0.07	0.14	0.56	1480
4	white	local & inc <49.5K	0.174	0.079	0.03	0.14	1.28	662
5	white	local & inc >49.5K	0.023	0.053	0.66	0.14	0.16	818
6	B-in-minority	all	0.038	0.020	0.05	0.17	0.22	4398
7	B-in-minority	arterial	0.029	0.028	0.30	0.17	0.17	2331
8	B-in-minority	local	0.048	0.025	0.06	0.18	0.27	2067
9	B-in-minority	local & inc <49.5K	0.078	0.035	0.03	0.20	0.39	1108
10	B-in-minority	local & inc >49.5K	0.012	0.038	0.76	0.16	0.07	959
11	mixed	all	0.036	0.015	0.02	0.18	0.20	2938
12	mixed	arterial	0.026	0.021	0.20	0.16	0.16	1638
13	mixed	local	0.047	0.025	0.06	0.22	0.22	1300
14	mixed	local & inc <49.5K	0.088	0.030	0.00	0.22	0.40	974
15	mixed	local & inc >49.5K	-0.073	0.046	0.11	0.21	-0.35	326
16	white + Black	all	0.021	0.014	0.12	0.16	0.13	3453
17	white + Black	arterial	0.020	0.015	0.19	0.15	0.13	2534
18	white + Black	local	0.022	0.030	0.47	0.19	0.11	919
19	white + Black	local & inc <49.5K	0.010	0.034	0.77	0.20	0.05	708
20	white + Black	local & inc >49.5K	0.065	0.058	0.26	0.17	0.38	211
21	Black	all	0.003	0.035	0.93	0.26	0.01	1245
22	Black	arterial	0.035	0.042	0.41	0.20	0.17	607
23	Black	local	-0.062	0.062	0.32	0.37	-0.17	638
24	Black	local & inc <49.5K	-0.047	0.063	0.46	0.35	-0.13	583
25	Black	local & inc >49.5K	-0.174	0.193	0.37	0.50	-0.35	55

The table reports the effects of the race of the pedestrian on the hit-and-run probability conditional on the block group type, road type, and the block group income of the crash location, using the GRF formula 7. Baseline means $E(h\&r|\text{white pedestrian})$.
Data: Fatal one-driver-one-pedestrian crashes in the U.S. in 2010-2016; non-Hispanic white and Black pedestrians only; sources: FARS, U.S. Census. **Abbreviations**: blgr – block group, inc – income, B-in-minority – *blacks-in-minority*.

Table A.12: Calibration test results: the test for heterogeneity in racial discrimination in hit-and-run rates across different types of block groups (race, income) and road types (arterial/local)

	Estimate	Std. Error	t value	p-val (est. ≤ 0)
mean grf prediction	1.018	0.241	4.23	0.0000
differential grf prediction	0.364	0.220	1.66	0.0487

Best linear fit using generalized random forest predictions of racial discrimination in hit-and-run rates as well as the mean forest prediction as regressors
Data: Fatal one-driver-one-pedestrian crashes in the U.S. in 2010-2016; non-Hispanic white and Black pedestrians only; sources: FARS, U.S. Census.

Table A.13: Probability of a hit-and-run in Texas: race of the pedestrian interacted with their residence; crashes that occurred in white, blacks-in-minority, or mixed block groups

			location fixed effects ℓ:	
	no controls (1)	all controls (2)	contig. clusters (3)	census tracts (4)
Black pedestrian	0.1052	0.0497	0.0435	0.0756
	(0.0207)	(0.0169)	(0.0175)	(0.0275)
	[0.0000]	[0.0035]	[0.0134]	[0.0061]
Resident	0.0130	0.0228	0.0165	0.0164
	(0.0123)	(0.0136)	(0.0145)	(0.0249)
	[0.2922]	[0.0931]	[0.2547]	[0.5110]
Black pedestrian × Resident	-0.0305	-0.0227	-0.0082	-0.0129
	(0.0224)	(0.0265)	(0.0272)	(0.0432)
	[0.1752]	[0.3910]	[0.7629]	[0.7646]
baseline⋆	0.1511	0.1665	0.1666	0.1573
N obs	3,930	3,930	3,930	3,962
N clusters			311	2,231
main controls	–	Yes	Yes	Yes
location FE:				
cont. cluster of tracts FE	–	–	Yes	–
census tract FE	–	–	–	Yes

⋆ Predicted probability of a hit-and-run, if everyone in the data were white pedestrians: $\mathbb{E}(\widehat{h\&r}|r = W, x)$;
This table reports the results for regression 2, where the indicator for the race of the pedestrian is interacted with the indicator for the residence status of the pedestrian. The pedestrian is considered to be a resident if the ZIP Code of their home address coincides with the ZIP Code of the crash location. Column (1) provides estimates without any controls. Column (2) reports the results with the full set of controls. Column (3) uses the contiguous racial clusters of census tracts for location fixed effects, while Column (4) uses census tracts. Standard errors – in parentheses – are clustered at the level of contiguous racial clusters of census tracts for columns (1), (2), and (3), and at the level of census tracts for column (4). P-values (H0: coef=0) are in brackets. **Data**: One-driver-one-pedestrian crashes in Texas with incapacitating injuries or fatalities, which occurred in 2010-2016 in block groups that are classified as white, blacks-in-minority, or mixed; non-Hispanic white and Black pedestrians only; sources: CRIS, U.S. Census.

Table A.14: Racial gaps in clearance rates by the type of the neighborhood and road. The GRF estimates.

	Blgr type	Sample subset	Estimate	S.E.	p-value	Baseline	Rel. gap	n obs
1	white	all	0.075	0.063	0.24	0.55	0.14	807
2	white	arterial	0.113	0.073	0.12	0.53	0.21	590
3	white	local	-0.016	0.111	0.88	0.61	-0.03	217
4	white	local & inc <49.5K	-0.136	0.154	0.38	0.64	-0.21	98
5	white	local & inc >49.5K	0.101	0.150	0.50	0.58	0.17	119
6	B-in-minority	all	-0.102	0.048	0.03	0.47	-0.22	812
7	B-in-minority	arterial	0.002	0.066	0.97	0.43	0.01	412
8	B-in-minority	local	-0.202	0.063	0.00	0.52	-0.39	400
9	B-in-minority	local & inc <49.5K	-0.271	0.067	0.00	0.51	-0.53	240
10	B-in-minority	local & inc >49.5K	-0.063	0.123	0.61	0.53	-0.12	160
11	mixed	all	-0.053	0.040	0.19	0.40	-0.13	627
12	mixed	arterial	0.028	0.059	0.64	0.37	0.08	287
13	mixed	local	-0.128	0.054	0.02	0.43	-0.30	340
14	mixed	local & inc <49.5K	-0.137	0.064	0.03	0.45	-0.30	274
15	mixed	local & inc >49.5K	-0.137	0.119	0.25	0.36	-0.38	66
16	white + Black	all	-0.025	0.042	0.55	0.50	-0.05	620
17	white + Black	arterial	-0.047	0.053	0.37	0.52	-0.09	421
18	white + Black	local	0.026	0.072	0.72	0.46	0.06	199
19	white + Black	local & inc <49.5K	0.009	0.084	0.91	0.44	0.02	155
20	white + Black	local & inc >49.5K	0.076	0.163	0.64	0.52	0.14	44
21	Black	all	-0.101	0.083	0.22	0.37	-0.28	368
22	Black	arterial	-0.093	0.097	0.34	0.38	-0.25	152
23	Black	local	-0.108	0.110	0.33	0.36	-0.30	216
24	Black	local & inc <49.5K	-0.062	0.122	0.61	0.33	-0.18	198
25	Black	local & inc >49.5K	-0.435	0.265	0.10	0.50	-0.87	18

The table reports the effects of the race of the pedestrian on the case clearance rate conditional on the block group type, road type, and the block group income of the crash location, using the GRF formula 7. Baseline means $E(identified|\text{white pedestrian}, h\&r = 1)$.
Data: Fatal one-driver-one-pedestrian hit-and-run cases in the U.S. in 2010-2016; non-Hispanic white and Black pedestrians only; sources: FARS, U.S. Census. **Abbreviations**: blgr – block group, inc – income, B-in-minority – *blacks-in-minority*.

Table A.15: Calibration test results: the test for heterogeneity in racial gap in the probability that the driver is identified across different types of block groups (race, income) and road types (arterial/local)

	Estimate	Std. Error	t-value	p-val (est ≤ 0)
mean GRF prediction	1.150	0.636	1.81	0.035
differential GRF prediction	0.377	0.230	1.64	0.051

Best linear fit using generalized random forest predictions of racial discrimination in hit-and-run rates as well as the mean forest prediction as regressors
Data: Fatal one-driver-one-pedestrian crashes in the U.S. in 2010-2016; non-Hispanic white and Black pedestrians only; sources: FARS, U.S. Census.

B (Online Appendix) Selection into the sample of fatal crashes and the direction of the bias

To understand the relationship between the estimated and underlying racial gap, consider the following numerical example. Suppose that every victim of a hit-and-run dies, while only half of the pedestrians die if the driver stays. If half of the drivers flee, independent of the race of the pedestrian, then conditional on a fatality, 66% of both Black and Black fatalities will be hit-and-runs, and the racial gap will be correctly estimated to be zero. If however, all the drivers flee when the pedestrian is Black, while only a half flees when the pedestrian is white, then conditional on a fatality, the share of hit-and-runs will be 100% for Black pedestrians and 66% for white pedestrians. Hence, we will estimate a 33 percentage point gap in hit-and-runs, which is lower than the true gap of 50 percentage points. We will also underestimate the relative racial gap in hit-and-runs: the 33 percentage point gap relative to the baseline 66% share of hit-and-runs for white fatalities implies a 50% increase in hit-and-runs, whereas in reality the drivers are 100% more likely to flee after hitting a Black pedestrian.

More formally, let us denote the mortality rate when the driver stays by $\omega_{0,x}$, and by $\omega_{0,x}$ when the driver flees. The mortality rate may differ depending on x: the circumstances of the crash, location, time, and age and sex of the pedestrian. Indeed, we may expect higher mortality rates in locations that are further away from hospitals, or for victims who are of old age, or for crashes on a highway, or for crashes that happen at night with no witnesses. However, conditional on all the crash circumstances, the expected share of hit-and-run fatalities out of all fatalities for a pedestrian of race r can be calculated by Bayes' rule as:

$$\mathbb{E}(h\&r|r; d, x, died) = \frac{\omega_{1,x} y_x^{d,r}}{\omega_{1,x} y_x^{d,r} + \omega_{0,x}(1 - y_x^{d,r})} \tag{B.1}$$

If drivers treat victims equally, i.e., $y_x^{d,r=B} = y_x^{d,r=W}$, then victims of both races of the same age, sex, and crash locations, should have the same expected share of hit-and-run fatalities out of all fatalities. If however drivers flee disproportionately more when the victim is (say) white, then we would also observe a higher share of hit-and-run fatalities out of all fatalities for white pedestrians in comparison to Black pedestrians. This is assuming that a hit-and-run may only aggravate the chances of the victim for survival, i.e., $\omega_1 \geq \omega_0$.

Proposition 1. If hit-and-runs cause more deaths, $\omega_{1,x} \geq \omega_{0,x}$, the expected gap in the hit-and-run rates for Black and white pedestrians has the same sign *conditional* or *unconditional* on the death of the pedestrian:

$$\text{sign}\left\{\mathbb{E}(h\&r|r = B; d, x, died) - \mathbb{E}(h\&r|r = W; d, x, died)\right\} = \text{sign}\left\{y_x^{d,B} - y_x^{d,W}\right\} \tag{B.2}$$

Proof. The formula for $\mathbb{E}(h\&r|r; d, x, died)$ in Eq. B.1 can be rewritten as

$$\left(1 + \frac{\omega_0}{\omega_1}\left(\frac{1}{y^r} - 1\right)\right)^{-1}$$

(subscripts x and superscripts d are omitted.) It is clear that $\mathbb{E}(h\&r|r; died)$ is monotonously strictly increasing in y^r. Hence if $y^B > y^W$, then $\mathbb{E}(h\&r|B; died) > \mathbb{E}(h\&r|W; died)$, and vice versa $\square$

Moreover, the percentage difference in hit-and-run rates for Black victims relative to white victims – i.e., the *relative* gap – will be always biased towards zero. In other words, if there is indeed discrimination of victims based on their race, then if we divide the absolute gap on the baseline probability, we get a conservative estimate of how much drivers more likely to flee when the pedestrian is discriminated against based on the skin color.

Proposition 2. As long as the mortality rate for a hit-and-run is greater or equal to the mortality for a hit-and-stay, $\omega_1 \geq \omega_0$, the relative gap in the hit-and-run rates for Black and white pedestrian fatalities will always underestimate the true percentage difference in the driver's propensity to flee for Black as opposed to white pedestrians:

$$\left|\frac{\mathbb{E}(h\&r|r = B; d, x, died) - \mathbb{E}(h\&r|r = W; d, x, died)}{\mathbb{E}(h\&r|r = W; d, x, died)}\right| \leq \left|\frac{y_x^{d,B} - y_x^{d,W}}{y_x^{d,W}}\right|$$

Proof.

$$\frac{\mathbb{E}(h\&r|d, BlackP = 1, died) - \mathbb{E}(h\&r|d, BlackP = 0, died)}{\mathbb{E}(h\&r|d, BlackP = 0, died)} = \frac{\nu y^{d,B} - y^{d,W}}{y^{d,W}}$$

where $\nu = \frac{\omega_s(1-y^{d,W})+\omega_r y^{d,W}}{\omega_s(1-y^{d,B})+\omega_r y^{d,B}}$. Then, we can show that $0 < \nu < 1$ if $y^{d,W} - y^{d,B} < 0$, while $\nu > 1$ if $y^{d,W} - y^{d,B} > 0$ □

To sum up, despite the likely selection bias inherent in the data, we can conclusively test whether there is discrimination in favor or against Black pedestrians. If the test reveals that there is discrimination, then the relative discrimination rate provides a lower bound to the true extent of the discrimination.

C (Online Appendix) Geo-spatial clustering based on racial characteristics

To cluster census tracts into contiguous areas based on the racial characteristics, we use the algorithm developed by Chodrow (2017b). The algorithm allows to regionalize spatial locations into segregation zones using information geometry. In particular, we use the R package *compx* by the same author. We follow closely the steps in Chodrow (2017a).

1. For a county c, prepare the map (shapefile) of census tract boundary and the data containing the vector $n^{(i)} = \{n_1^{(i)}, \ldots, n_5^{(i)}\}$ with the count of residents per each ethnic group: i.e., $n_1^{(i)}$ is a number of non-Hispanic white residents: $n_2^{(i)}$, Hispanic white; $n_3^{(i)}$, Black; $n_4^{(i)}$, Asian; and $n_5^{(i)}$, other residents. The shapefiles and the data are based on census data for 2010, retrieved from Manson et al. (2017).

2. Compute graph $g^{(c)}$ using function `construct_information_graph()`. and the Jensen-Shannon metric to calculate pairwise divergence between any two adjacent tracts i and j, where $i, j \in tracts(c)$:

$$\sqrt{\frac{N^{(i)}}{N^{(i)} + N^{(j)}} D_{KL}(p^{(i)}, p^{(i)} + p^{(j)}) + \frac{N^{(j)}}{N^{(i)} + N^{(j)}} D_{KL}(p^{(j)}, p^{(i)} + p^{(j)})}$$

   where $N^{(k)} = \sum_{l=1}^{5} n_l^{(k)}$ is a total population in census tract k; $p^{(j)} = \{n_1^{(k)}, \ldots, n_5^{(i)}\}/N^{(k)}$ is a vector with a share of each ethnic group in census tract k; and $D_{KL}(.,.)$ is the Kullback-Leibler divergence (relative entropy) $D_{KL}(p^{(i)}, p^{(i)} + p^{(j)}) = \sum_{l=1}^{5} p_l^{(i)} \log\left(\frac{p_l^{(i)}}{p_l^{(i)} + p_l^{(j)}}\right)$.

3. Compute the Laplacian matrix corresponding to the graph $g^{(c)}$.

4. Request 100 eigenvalues of the the Laplacian matrix (or as many eigenvalues as the number of rows of the Laplacian matrix, whicher is more binding).

5. Select the optimal number of clusters k by finding at which position from the end the eigenvalue becomes larger than 0.01.[19]

6. Perform k-means clustering in the eigenspace of the Laplacian matrix 1000 times and pick the one with the best performance. Select another county and repeat steps 1 to 6.

D (Online Appendix) Checking for omitted interactions with double lasso

The double-lasso method incorporates lasso's model selection properties into the causal inference framework. Lasso regression is a penalized regression model that performs model selection for prediction tasks by shrinking some parameters to exactly zero (Tibshirani 1996). In particular, it transforms a high-dimensional inference problem into a low-dimensional inference problem, by partialling out nuisance variables from the outcome variable *hit-and-run*, and from the variable of interest *Black pedestrian* using the variables that have been selected by two separate lasso regressions.

In our case, the high-dimensional inference problem is $h\&r_i = \beta I\{r = B\}_i + f(X_i) + e_i$, where $f(X)$ is a potentially non-linear transformation of nui-

19. "[T]here is some judgment required in identifying the number of" clusters that "most fully describes the community structure." See *The scale of segregation* section in Chodrow (2017a). The author of the algorithm suggests finding the number of clusters by visually inspecting the gaps in the eigenvalues, and picking the one out of several potential candidates. Since in our paper, we had to cluster automatically for 2,036 counties, we had to use a simplified data-driven approach with a fixed cut-off, rather than relying on a visual inspection.

sance variables X. The low-dimensional representation is $h\&r_i - \widehat{h\&r}(X_i)) = \beta(I\{r = B\}_i - I\widehat{\{r = B\}}(X_i)) + u_i$, where $\widehat{h\&r}(X_i)$ and $I\widehat{\{r = B\}}(X_i)$ are the predictions based on the nuisance variables X and their interaction terms. To use the double-lasso method, we assume approximate sparsity, i.e., only some interaction terms (not all of them) can be significant confounding factors, simultaneously explaining both the probability that the pedestrian is Black and the probability that the driver runs away.

We create all possible pairwise interaction terms between the controls in our main specification – the set of block group, crash, pedestrian characteristics and all the time dummies – and end up with 5,502 non-collinear potential controls, including the square terms. In order to avoid creating many dummy variables with little variation, we do not interact the time trends separately for weekends with other variables, but we do interact time trends with other variables in general. Also, we treat the number of lanes as continuous variable only for interaction terms, keeping the original dummies for number of lanes as controls. Next, we partial out the fixed effects of the geographical clusters from the outcome variable *hit-and-run*, the variable of interest *Black pedestrian*, and all other 5,502 potential controls.

Then, we use rigorous lasso (Chernozhukov et al. 2016; Belloni et al. 2012) to predict hit-and-runs, which selects 29 candidate interaction terms that could predict variation in hit-and-runs:

- *blgr_white_s*×*VNUM_LAN*;
- *blgr_white_s*×*AGE_P_sqrd*;
- *blgr_asian_s*×*HOUR14*;
- *POP_urban_s*×*HOUR2*;
- *POP_urban_s*×*HOUR3*;
- *POP_urban_s*×*HOUR14*;
- *POP_urban_s*×*LGT_CONDDark – Lighted*;

- *monthyear32×HOUR10*;
- *monthyear41×HOUR11*;
- *monthyear43×HOUR11*;
- *monthyear24×HOUR12*;
- *monthyear43×HOUR15*;
- *monthyear5×HOUR16*;
- *monthyear17×HOUR20*;
- *monthyear15×HOUR22*;
- *monthyear35×DAY_WEEK4*;
- *monthyear62×DAY_WEEK6*;
- *monthyear29×LGT_CONDDark – Unknown Lighting*;
- *HOUR9×DAY_WEEK3*;
- *HOUR11×DAY_WEEK3*;
- *HOUR15×arterial_crash*;
- *HOUR18×arterial_crash*;
- *HOUR1×AGE_P*;
- *HOUR2×AGE_P*;
- *HOUR3×AGE_P*;
- *HOUR11×AGE_P_sqrd*;
- *HOUR19×intoxicated_pedestrian*;
- *weekend×LGT_CONDDark – Lighted*;
- *VNUM_LAN×AGE_P_sqrd*

We separately apply the same procedure to predict the race of the pedestrian, for which rigorous lasso chooses three additional interaction terms:

- blgr_white_s×AGE_P;
- POP_urban_s×AGE_P;
- monthyear54×HOUR1

Finally, we re-run Regression 2, combining both sets of interaction terms chosen by lasso with our initial set of controls. As shown in Section3, including these terms does not affect the result.

E (Online Appendix) Exploring potential mechanisms: a theoretical framework

Consider a simple model based on Becker (1974). A driver i chooses whether to run away after hitting a pedestrian of race r, following a crash with characteristics x (e.g., location, time, other salient characteristics of the victim), if his expected utility from staying is less than the expected utility from fleeing, i.e.:

$$h\&r_x^{i,r} = \begin{cases} 1, & \text{if } \pi_x^{i,r} - \omega_x^r \leq -p_x^r(\omega_x^r + \Omega_x^r) \\ 0, & \text{otherwise} \end{cases} \tag{E.3}$$

where $\pi_x^{i,r} \in [-\infty, +\infty]$ is a subjective net psychological utility from staying, ω_x^r is the expected legal sanction for hitting the pedestrian, p_x^r is the probability of being caught, and Ω_x^r expresses the expected legal sanctions for failing to stop and render aid, i.e., the hit-and-run penalty.

The driver runs away if his subjective net psychological utility from staying is larger than the threshold based on legal sanctions, τ_x^r:

$$h\&r_x^{i,r} = 1 \text{ iff } \pi_x^r \leq \tau_x^r \equiv (1 - p_x^r)\omega_x^r - p_x^r\Omega_x^r \tag{E.4}$$

For example, very empathetic and law-abiding drivers have large and positive π and would thus be more likely to stay. Drivers that are in a state of shock (as some would say, those that are acting irrationally) can be seen as having large negative π. Such drivers would flee independent of the circumstances, even in front of multiple witnesses. Similarly, drivers who believe they are more at fault in the accident expect higher legal sanctions ω if they stay, so such drivers have higher incentives to flee to avoid punishment. Indeed, as we see from the descriptive statistics, identified hit-and-run drivers had aggravating circumstances, such as driving without a license, or under intoxication. Importantly, causal evidence from Castriota and Tonin (2019) shows that drivers are more likely to flee in circumstances where they expect a lower probability of being identified later – such as under the veil of darkness.

Assuming that the net subjective psychological utility from staying is distributed among the drivers at risk following $F_x^r(.)$, then the probability of a hit-and-run given the race of the victim r and the characteristics of the crash x is:

$$E(h\&r_x^{i,r}) = F_x^r(\tau_x^r) \tag{E.5}$$

In the data, we find that drivers react differently to the race of the pedestrian, i.e., $F_x^B(\tau_x^B) > F_x^W(\tau_x^W)$. This can be explained by either (1) $\tau_x^B > \tau_x^W$ – i.e., drivers expect that the legal sanctions are lower when the victim is Black – or (2) by $F_x^B(\tau) > F_x^W(\tau)$ – i.e., drivers expect the same sanctions but attach lower utility to helping an injured Black pedestrian on average, or (3) both. The first case would be an instance of *statistical discrimination*, and the second would be seen as *taste-based discrimination*: for example, the result of in-group/out-group bias.